IIAR JOURNALS

Innovating Knowledge. Advancing Research.

IIAR Journal of Security, Peace, and Conflict Studies

Vol 1, No 1, 2026

Artificial Intelligence and the Changing Character of Armed Conflict: A Theoretical Framework of Autonomy, Speed and Accountability in Military Decision-Making

NWOGBOJI Emmanuel Usulor[1], OGAH SUNDAY OGBO[2], Atolagbe Umar Kawu[3], Paul Idowu Bamigbola[4]

[1,2,3,4] Veristzon Global Research Lab

ABSTRACT

Artificial intelligence is entering warfare not only as autonomous weapons but, more consequentially, as decision support that compresses the time between detection and strike, selects targets from data at scale, and shapes what commanders see. This theoretical paper develops a framework for the effects of AI on military decision-making. Drawing on strategic studies and the theory of war, on the international-humanitarian-law and meaningful-human-control debates, on the civil–military-relations literature, and on recent research on AI-enabled targeting, decision support and autonomy in contemporary conflicts, it distinguishes three changes AI introduces, autonomy (decisions without a human in the loop), speed (decisions faster than human deliberation) and opacity (decisions whose basis cannot be reconstructed), and argues that the second and third matter more than the first, because a human nominally in the loop of a system that recommends at machine speed from unexplainable inference exercises no meaningful control. The paper derives seven propositions about the effects of AI on escalation dynamics, on civilian harm and on the attribution of responsibility for unlawful acts, and considers the position of African states and regional organisations that will encounter these systems as importers, as hosts of conflicts in which others deploy them, and as parties to the international negotiations on their regulation. Nigeria's use of armed drones in the Sahel and Lake Chad Basin anchors the African analysis. The paper closes with implications for military doctrine, arms-control diplomacy and the research agenda of African security studies. It concludes that a human in the loop is not in itself a safeguard but a question, since whether that human can deliberate, whether she has the time, the reasons and the authority, is what determines whether the loop is hers or the machine's.

Keywords: artificial intelligence; armed conflict; autonomous weapons; military decision- making; meaningful human control

Artificial Intelligence and the Changing Character of Armed Conflict: A Theoretical Framework of Autonomy, Speed and Accountability in Military Decision-Making

1. INTRODUCTION

The debate on artificial intelligence in war has been organised, for a decade, around the lethal autonomous weapon: the machine that selects and engages targets without a human decision. That debate has produced a large literature, a United Nations process and a vocabulary, meaningful human control, human in the loop, that now frames how states and publics think about the problem (Bode et al., 2023; Qerimi, 2023; Geith, 2026). It has also, this paper argues, directed attention to the wrong place. The systems that have actually changed the character of recent armed conflict are not autonomous weapons in the strict sense. They are decision-support systems that process intelligence at scale, generate target lists, estimate the location and status of individuals, and present recommendations to human operators who approve them in seconds (Andersin, 2025; Wolf, 2025; Abbasi & He, 2026). The human remains in the loop. What has changed is the loop.

This paper develops a theoretical framework for the effects of AI on military decision-making that takes this observation as its starting point. It distinguishes three changes that AI introduces, autonomy, speed and opacity, and argues that the second and third are more consequential than the first, because they operate even where, and especially where, a human remains formally responsible. A commander who approves a hundred algorithmically generated targets in an hour, each on the basis of a probability she cannot interrogate, is in the loop in the sense the autonomous-weapons debate recognises and out of it in every sense that matters. The paper's framework is built to capture that distinction.

The framework draws on four literatures: strategic studies and the theory of war, which supply the concepts of escalation, tempo and strategic stability (Johnson, 2023, 2026a, 2026b; Horowitz, 2026; Edmond, 2026); the international-humanitarian-law scholarship on targeting, precautions and proportionality as applied to AI decision support (Andersin, 2025; Wolf, 2025; Abbasi & He, 2026); the meaningful-human-control and human–machine-teaming literatures (Bode et al., 2023; Lushenko, 2025; Tiwari & Babu, 2026); and the emerging research on military AI in Africa (Okpaleke et al., 2023; Oyewole et al., 2025). From these it derives propositions about escalation, civilian harm and accountability, and it situates African states in the resulting landscape.

The analysis begins with the autonomous-weapons debate and why it has mis-specified the problem, and then locates where AI actually enters the targeting cycle. It develops the three-change decomposition of autonomy, speed and opacity before reconsidering the human-in-the-loop standard against it and deriving propositions on escalation, civilian harm and accountability. It closes by situating African states in the resulting landscape and drawing out the implications for doctrine, law and arms control.

2. LITERATURE REVIEW

2.1 The Autonomous-Weapons Debate and Its Blind Spot

The debate

Since the early 2010s the governance of military AI has been organised around lethal autonomous weapon systems and the question of whether, and how, to prohibit or regulate weapons that select and engage targets without human intervention. Bode et al. (2023) compare Chinese, Russian and United States practices and find that the prospects for global governance are shaped less by shared principle than by divergent operational practice, with major powers developing autonomy in ways that outrun the negotiations. Qerimi (2023) offers a typology of state positions ranging from prohibition to permissive regulation, and Geith (2026) explains the multilateral bargaining as a lowest-common-denominator process in which the most permissive states set the ceiling of what can be agreed. Johansen and Kruck (2025) add a political-economy dimension: states face a competence–control trade-off in which the private experts who build autonomous systems acquire influence that erodes state control over military innovation. The debate has been productive in establishing that autonomy raises problems of legal compliance, accountability and strategic stability that existing arrangements do not address.

The blind spot

The debate's organising assumption, that the critical line is between human and machine selection of targets, has a blind spot that recent conflicts have exposed. The systems that have transformed targeting in practice keep a human at the point of decision while transforming everything before it.

© IIARJournal. All rights reserved.

Andersin (2025), analysing the reported use of the Lavender system in Gaza, shows that an AI decision-support system combined with facial-recognition technology can generate target identifications at a scale and speed that the law of targeting's requirements of verification and precaution were not designed to meet, with the human approval step reduced to a brief check. Wolf (2025) examines the obligation to take precautions in target verification when the verification itself is performed by a model, and finds that the human operator's capacity to discharge the obligation depends on information the system may not provide. Abbasi and He (2026) argue that AI decision support reconfigures proportionality, an inherently qualitative judgement under international humanitarian law, by presenting it as a quantitative output, and that the reconfiguration shifts the locus of judgement away from the human who is formally responsible for it. Vanderborght and Nadibaidze (2025) show, from a different angle, that the virtual demonstrations through which such systems are presented shape perceptions of war and security before any deployment, normalising speed and scale as virtues.

The blind spot, then, is this: a governance framework organised around autonomy will certify as compliant a system in which a human approves every strike while being unable to deliberate on any. The following section proposes a decomposition that brings the real changes into view.

2.2 Where AI Enters the Targeting Cycle

To see why speed and opacity matter more than autonomy, it helps to locate where artificial intelligence actually enters military decision-making. The targeting cycle, in the doctrine of most modern militaries, proceeds through stages: the collection of intelligence; its processing and fusion into a picture of the battlespace; the development of targets from that picture; the validation of targets against legal and operational criteria; the engagement decision; and the assessment of effects afterwards. The autonomous-weapons debate has concentrated on the engagement decision, the point at which a weapon is released. AI has entered earlier and more pervasively.

Collection and processing

The first and least contested entry point is the processing of intelligence. Modern collection, satellite and aerial imagery, signals, communications metadata, open sources, biometric and facial data, produces volumes no human analytical staff can read. Machine learning sorts, classifies, links and prioritises, and in doing so decides what the human will see. Vanderborght and Nadibaidze (2025) show that the systems performing this work are presented, in military demonstrations, as neutral accelerators of situational awareness, when in fact they shape the picture of the battlespace from which every later decision follows. An entity that the model has flagged becomes a candidate; one it has not flagged does not exist for the purposes of the cycle. Opacity enters here, before any targeting decision is made.

Target development and validation

The second entry point is target development: the generation of lists of persons, locations and objects nominated for engagement. This is where the systems analysed by Andersin (2025) and Wolf (2025) operate. A model trained on data about known members of an adversary organisation scores the population for similarity, and those above a threshold become targets for validation. The scale of generation is bounded only by computation; the scale of validation is bounded by the number of trained humans and the time each can give. Speed enters here, as the ratio between the two. Validation, the stage at which international humanitarian law's requirements of distinction, precaution and proportionality are meant to be discharged, becomes the bottleneck, and the pressure to clear the bottleneck falls on the human who must approve or reject each nomination. Abbasi and He (2026) show that proportionality, in such a system, is increasingly presented to the validator as an output rather than left to her as a judgement.

Engagement and assessment

The engagement decision is where autonomy would enter, and where, in most contemporary systems, it has not. The human presses the button. But a human pressing a button at the end of a cycle whose earlier stages have been automated is approving a product she did not make, under time pressure the earlier stages created, on the basis of reasons the earlier stages did not expose. Assessment, the final stage, is similarly transformed: battle-damage assessment performed by the same models that generated the target will tend to confirm the target, and the feedback loop that should correct error instead reinforces it.

The lesson

© IIARJournal. All rights reserved.

The lesson of the cycle is that AI's effect on military decision-making is cumulative and upstream. By the time the engagement decision arrives, the human's capacity to deliberate has been determined by what the earlier stages selected, how fast they delivered it and whether they explained themselves. A governance framework that inspects only the engagement decision inspects the last and least consequential link. The three changes examined next are properties of the whole cycle.

3. CONCEPTUAL FRAMEWORK

3.1 Three Changes: Autonomy, Speed and Opacity

AI introduces three analytically distinct changes into military decision-making. Each can occur without the others, and each has its own consequences for law, strategy and accountability.

Autonomy

Autonomy is the removal of the human from the decision: a system that selects and engages targets, or takes other consequential actions, without a human approval. It is the change the governance debate has addressed, and its problems are well specified. Autonomy breaks the chain of individual responsibility that international humanitarian law and criminal law presuppose; it removes the moral agent from the act; and it raises the prospect of systems interacting with one another at a tempo and in ways no human intended (Bode et al., 2023; Johnson, 2026a). It is also, so far, the least common change in practice. Fully autonomous engagement remains confined to defensive systems against incoming munitions and to a small number of loitering munitions; the great majority of AI in contemporary targeting is decision support.

Speed

Speed is the compression of the decision cycle below the time required for human deliberation. Johnson (2023), examining the automation of the observe–orient–decide–act loop, argues that the acceleration of command-and-control decision-making by AI threatens to remove the human from the parts of the loop that cannot be delegated, judgement, context, moral evaluation, while retaining a human presence that is nominal rather than deliberative. Speed does not require autonomy. A system that presents a human with a target and a ten-second window, in a queue of hundreds, has not removed the human; it has removed the deliberation. Tiwari and Babu (2026), drawing lessons from human-in-the-loop combat AI systems in Ukraine, document human–machine teaming under conditions in which the human's contribution is constrained by the tempo the machine sets. Speed has strategic as well as tactical consequences: Horowitz (2026) and Edmond (2026) examine how AI-driven acceleration affects strategic and crisis stability, and Johnson (2026b) revisits the stability–instability paradox under AI-enabled warfare, arguing that the combination of conventional AI advantage and nuclear shadow creates escalation risks that classical deterrence theory did not anticipate. Ali and Zou (2026) apply the same concern to South Asia, where AI-enabled command systems could inadvertently escalate crises between nuclear-armed rivals.

Opacity

Opacity is the inability to reconstruct the basis of a decision. A machine-learning model that identifies a target from patterns in data does not provide reasons in the sense that a human analyst does; it provides a score. The operator who approves the target cannot know which features drove the score, whether the model's training reflects the current situation, or how confident the score should make her. Wolf (2025) shows that this undermines the precautionary obligation to verify, because verification requires information the system does not expose. Abbasi and He (2026) show that it undermines proportionality assessment, because the qualitative weighing the law requires cannot be performed on a quantitative output whose derivation is hidden. Johnson (2026a), asking whether AI can behave ethically in military crises, identifies automation bias, anthropomorphism and over-reliance as the mechanisms by which opacity translates into the surrender of human moral agency. Lushenko (2025), in a conjoint experiment in the United States military, finds that trust in AI varies with the level of decision and with institutional endorsement, suggesting that the conditions under which operators defer to opaque systems are themselves shaped by how the systems are introduced. Opacity, like speed, does not require autonomy. It is a property of the inference, not of the decision rule.

The interaction

The three changes interact, and the interaction is where the character of conflict changes. Speed without opacity, a fast system whose reasoning is transparent, could in principle be checked by a well-resourced human. Opacity without speed, a slow, unexplainable system, could be interrogated over time. Speed and opacity together produce a system

© IIARJournal. All rights reserved.

in which the human cannot deliberate because there is no time and could not deliberate because there is nothing to deliberate on. Autonomy, in this light, is less a separate category than the limit case: the point at which the human's contribution has become so attenuated that removing it changes nothing. The governance debate has guarded the limit case while the interaction has transformed practice.

3.2 The Human-in-the-Loop Standard Reconsidered

What the standard requires

The human-in-the-loop standard, and its more demanding formulation as meaningful human control, holds that a human must make or approve consequential decisions and must do so with sufficient information, time and authority to make the decision her own. Bode et al. (2023) and the broader literature treat this as the condition that preserves legal responsibility and moral agency. The standard is sound. The problem is its application.

What the standard certifies

Applied to autonomy alone, the standard certifies any system in which a human approves the output. Applied to the three changes together, it would require that the human have time to deliberate (a constraint on speed) and reasons to deliberate on (a constraint on opacity). The autonomous-weapons debate has applied the first version. The consequence, visible in the targeting systems analysed by Andersin (2025), Wolf (2025) and Abbasi and He (2026), is that systems which satisfy the letter of human control have transformed the practice of targeting beyond what the standard was meant to permit. The human is in the loop; the loop runs at machine speed on machine inference; and the human's approval is a formality that transfers legal responsibility to her without transferring the capacity to exercise it.

Meaningful control as a property of the system, not the decision rule

The framework proposed here relocates meaningful human control from the decision rule, is there a human approval step?, to the system as a whole: does the system, in its tempo, its information outputs and its institutional setting, allow the human to deliberate? This yields three tests. The time test asks whether the interval between recommendation and required decision permits the human to consider the recommendation in context. The reasons test asks whether the system exposes the basis of its recommendation in a form the human can evaluate against the legal and operational standards she is responsible for. The authority test asks whether the human can reject the recommendation without institutional penalty and whether rejection is in practice exercised. A system that fails any test does not afford meaningful control whatever its decision rule. Tiwari and Babu (2026) and Lushenko (2025) supply evidence that the authority test, in particular, is frequently failed: operators under tempo pressure and institutional endorsement defer.

Consequences for law

The relocation has legal consequences. If responsibility attaches to the human who approves, but the system does not afford meaningful control, then responsibility attaches to someone who could not have exercised it, and accountability becomes a fiction that protects the system's designers and the commanders who deployed it. The responsibility-gap literature has treated this as a problem of autonomy; the framework shows it to be a problem of speed and opacity as well, and therefore far more widespread. Abbasi and He (2026) and Wolf (2025) point toward the remedy: obligations of precaution and proportionality must be read as obligations on the system's design and deployment, not only on the operator's approval.

3.3 Propositions

The framework generates propositions about three outcomes: escalation, civilian harm and accountability.

Escalation

Proposition 1. The escalation risk of AI in military decision-making rises with speed and opacity independently of autonomy, because decision cycles compressed below the threshold of deliberation remove the pauses in which de-escalation occurs, and opaque systems deprive decision-makers of the information needed to recognise error before it compounds (Johnson, 2023, 2026b; Horowitz, 2026). Proposition 2. Where adversaries both deploy AI-enabled decision support, the interaction of their systems will produce escalation dynamics that neither side intended and neither side can reconstruct, and the risk is highest between nuclear-armed rivals in conventional confrontation (Johnson, 2026b; Ali & Zou, 2026; Edmond, 2026).

Civilian harm

Proposition 3. AI decision support increases civilian harm where it increases the number of targets generated per unit of verification capacity, because

© IIARJournal. All rights reserved.

the obligations of verification and precaution are discharged by humans whose time does not scale with the system's output (Andersin, 2025; Wolf, 2025). Proposition 4. Quantitative presentation of proportionality by AI systems will shift the threshold of acceptable harm, because a numerical output invites comparison against a numerical standard that the qualitative legal test does not supply, and operators under tempo pressure will treat the output as the assessment rather than as an input to it (Abbasi & He, 2026). Proposition 5. The civilian-harm effect of AI decision support is greatest in conflicts where the adversary is embedded in the civilian population and where intelligence is drawn from population-wide data, because the model's identifications rest on patterns of association rather than on observed conduct, and the error rate is borne by civilians.

Accountability

Proposition 6. Responsibility gaps arise wherever the time, reasons or authority tests are failed, not only where autonomy is present; accountability frameworks organised around autonomy will therefore miss most of the cases in which no one can be held responsible for an unlawful act. Proposition 7. Attribution of responsibility for AI-enabled harm will migrate from operators to designers, procurers and commanders only where law and doctrine impose obligations on system design and deployment; absent such obligations, the operator who approved will be held responsible for a decision she could not have made differently.

The propositions are stated so as to be testable through comparative study of conflicts in which AI decision support has been deployed, through experimental work on operator decision-making under tempo and opacity, and through doctrinal analysis of accountability outcomes.

4. ANALYSIS AND DISCUSSION

4.1 African States in the Landscape of Military AI

African states encounter military AI in three positions, and the framework has different implications for each.

Importers

African militaries do not, for the most part, develop AI-enabled systems; they acquire them. Oyewole et al. (2025), examining autonomous weapons systems in Africa, document emerging realities, prospects and risks across the continent and note that the African Union and a substantial number of African states have taken positions in the international process on autonomous weapons, while acquisition proceeds in advance of regulation. Okpaleke et al. (2023), making the case for drones in counter-insurgency in the West African Sahel, show that Nigeria and other states have employed armed drones against insurgents and that the use has outpaced the development of doctrine, legal review and civilian-protection practice. As importers, African states acquire systems whose autonomy, speed and opacity are designed elsewhere and whose meaningful-control properties they cannot modify. The time, reasons and authority tests must therefore be applied at procurement: a state that cannot establish whether a system affords deliberation should not deploy it in targeting. Nigeria's armed-drone operations in the north-east and north-west illustrate the stakes. Targeting against groups embedded in rural populations is the condition Proposition 5 identifies as most hazardous, and the addition of AI-enabled target generation to such operations would multiply the risk unless verification capacity scaled with it.

Hosts

African states are also hosts of conflicts in which external actors deploy AI-enabled systems: counter-terrorism operations by foreign militaries, private military companies and partner forces operating with AI-enabled intelligence and targeting support. In this position the state neither controls the system nor, often, knows its properties, while its population bears the civilian-harm effect. The framework implies that host states should require, as a condition of consent to foreign operations, disclosure of the decision-support systems in use and of the verification arrangements that accompany them, and should treat failure of the time and reasons tests as grounds for withholding consent.

Negotiators

African states are parties to the international negotiations on autonomous weapons, and Oyewole et al. (2025) note the African Union's engagement. The framework suggests that the African position, which has generally favoured prohibition of fully autonomous systems, is well aimed at the limit case and poorly aimed at the practice. A negotiating position that extended meaningful human control to speed and opacity, requiring that decision-support systems afford deliberation, would address the systems that will be used in and against African states, and would align with the interests of states

© IIARJournal. All rights reserved.

that are importers and hosts rather than developers. Geith's (2026) account of lowest-common-denominator bargaining suggests that such a position would face resistance from developer states, but Qerimi's (2023) typology shows that coalitions of non-developer states have shaped the terms of the debate before.

The verification constraint

Across all three positions, the binding constraint for African states is verification capacity. The propositions locate civilian harm and accountability failure in the ratio between the rate at which systems generate targets and the capacity of trained humans to validate them. African militaries, with smaller intelligence staffs, thinner legal-advisory capacity and less developed targeting doctrine than the developer states, have the lowest verification capacity and would therefore experience the sharpest degradation of meaningful control from any given system. This suggests a counter-intuitive conclusion: the states with least capacity to absorb AI decision support are the states for which it is most dangerous, and the appropriate response is not to acquire the systems more slowly but to invest in verification, trained targeting officers, legal advisers, assessment capability, before and as a condition of acquisition. Oyewole et al. (2025) observe that acquisition across the continent has run ahead of regulation; the framework adds that it has run ahead of the human capacity that regulation presupposes.

4.2 Implications

For doctrine

Military doctrine on AI-enabled decision support should incorporate the time, reasons and authority tests as operational requirements: minimum deliberation intervals for consequential decisions, mandatory exposure of the basis of recommendations in evaluable form, and institutional protection for operators who reject recommendations. Doctrine should treat verification capacity as a constraint on the rate of target generation rather than allowing generation to set the rate of verification.

For law

International humanitarian law should be read, and where necessary developed, to impose obligations of precaution and proportionality on the design and deployment of decision-support systems, not only on the operator's approval. Accountability frameworks should recognise that responsibility gaps arise from speed and opacity as well as from autonomy, and should allocate responsibility to those who designed, procured and deployed systems that did not afford meaningful control.

For arms control

The international process should extend its scope from autonomous weapons to AI-enabled decision support, and should develop standards for meaningful human control that address tempo and explainability. African states, as importers, hosts and negotiators, have a shared interest in such an extension and the collective weight to advance it.

For research

The propositions can be tested. The civilian-harm effects of AI target generation can be studied through comparison of incidents across operations with and without such systems. Operator decision-making under tempo and opacity can be studied experimentally, extending Lushenko (2025). The escalation dynamics of interacting AI systems can be modelled. And African cases, Nigeria's drone operations above all, need the peer-reviewed study they have largely lacked, on their own terms and with attention to the civilian-protection questions the framework raises.

Limits

The paper is conceptual and relies on a small number of documented cases. It treats the three changes as separable when deployed systems combine them. It draws its evidence on contemporary targeting from a single theatre whose documentation is contested. And it does not address AI in intelligence, logistics, cyber or information operations, where the three changes also apply.

5. CONCLUSION AND RECOMMENDATION

The governance of military artificial intelligence has been organised around a line between human and machine decision that the technology has largely left behind. The systems transforming contemporary conflict keep a human at the point of decision and transform everything before it: they generate targets at scale, present recommendations at speed, and rest on inference no human can reconstruct. This paper has decomposed the changes AI introduces into autonomy, speed and opacity; argued that the second and third are the more consequential because they operate through, rather than around, the human in the loop; relocated meaningful human control from the decision rule to the system; derived propositions about escalation, civilian harm and

© IIARJournal. All rights reserved.

accountability; and examined the position of African states as importers, hosts and negotiators.

The conclusion is that a human in the loop is not a safeguard but a question. Whether she can deliberate, whether she has the time, the reasons and the authority, is what determines whether the loop is hers or the machine's. States that acquire, host or negotiate the regulation of these systems, Nigeria and its African neighbours among them, should ask that question of every system, and should decline to deploy, consent to or permit those that cannot answer it.

6. CONTRIBUTION TO EXISTING KNOWLEDGE

The paper contributes a decomposition of the changes artificial intelligence introduces into military decision-making into autonomy, speed and opacity, and argues that the second and third matter more than the first because they operate through, rather than around, the human who remains formally in the loop. It relocates the standard of meaningful human control from the decision rule to the system as a whole, yielding tests of time, reasons and authority that a decision-support system must satisfy.

By showing that responsibility gaps arise from speed and opacity as well as from autonomy, the paper extends accountability analysis beyond the autonomous-weapon case that dominates the governance debate, and it sets out the implications for military doctrine, international humanitarian law and arms-control diplomacy, with particular attention to the position of African states as importers, hosts and negotiators.

7. REFERENCES

Abbasi, A. Y., & He, J. (2026). AI-enabled decision support systems and the reconfiguration of proportionality in international humanitarian law. Computer Law & Security Review, 61, 106319. https://doi.org/10.1016/j.clsr.2026.106319

Ali, J., & Zou, X. (2026). AI security dilemma and militarisation in South Asia: Rethinking strategic stability in the age of autonomous warfare. Sage Open, 16(3), 21582440261475709. https://doi.org/10.1177/21582440261475709

Andersin, E. (2025). The use of the 'Lavender' in Gaza and the law of targeting: AI-decision support systems and facial recognition technology. Journal of International Humanitarian Legal Studies, 16(2), 336–370. https://doi.org/10.1163/18781527-bja10119

Bode, I., Huelss, H., Nadibaidze, A., Qiao-Franco, G., & Watts, T. F. A. (2023). Prospects for the global governance of autonomous weapons: Comparing Chinese, Russian, and US practices. Ethics and Information Technology, 25(1), 5. https://doi.org/10.1007/s10676-023-09678-x

Edmond, C. (2026). Autonomous weapons and strategic stability. Discover Global Society, 4(1), 221. https://doi.org/10.1007/s44282-026-00548-7

Geith, J. (2026). Lowest common denominator: Explaining multilateral bargaining over lethal autonomous weapons systems regulation. Global Policy, 17(1), 13–35. https://doi.org/10.1111/1758-5899.70101

Horowitz, M. C. (2026). Artificial intelligence and the future of strategic stability. Texas National Security Review, 9(2), 66–83. https://doi.org/10.1353/tns.00034

Johansen, A., & Kruck, A. (2025). The competence–control trade-off in military AI innovation: Autonomous weapons systems and shifting modes of state control over private experts. European Journal of International Security. Advance online publication. https://doi.org/10.1017/eis.2025.10020

Johnson, J. (2023). Automating the OODA loop in the age of intelligent machines: Reaffirming the role of humans in command-and-control decision-making in the digital age. Defence Studies, 23(1), 43–67. https://doi.org/10.1080/14702436.2022.2102486

Johnson, J. (2026a). Can AI behave ethically during military crises? Preserving human moral agency. International Affairs, 102(1), 63–83. https://doi.org/10.1093/ia/iiaf191

Johnson, J. (2026b). Revisiting the 'stability–instability paradox' in AI-enabled warfare: A modern-day Promethean tragedy under the nuclear shadow? Review of International Studies, 52(5), 1050–1068. https://doi.org/10.1017/s0260210524000767

Lushenko, P. (2025). AI, trust, and the war room: Evidence from a conjoint experiment in the US military. Cambridge Forum on AI: Law and Governance, 1, e52. https://doi.org/10.1017/cfl.2025.10019

Okpaleke, F. N., Nwosu, B. U., Okoli, C. R., & Olumba, E. E. (2023). The case for drones in counter-insurgency operations in West African Sahel. African Security Review, 32(4), 351–367. https://doi.org/10.1080/10246029.2023.2217158

Oyewole, S., Isike, C., Oche, T., & Olumba, E. E. (2025). Autonomous weapons systems in Africa: Emerging realities, prospects, and risks. Journal of Applied Security Research, 20(4), 603–628. https://doi.org/10.1080/19361610.2025.2501055

Qerimi, Q. (2023). Controlling lethal autonomous weapons systems: A typology of the position of

© IIARJournal. All rights reserved.

states. Computer Law & Security Review, 50, 105854. https://doi.org/10.1016/j.clsr.2023.105854

Tiwari, R. K., & Babu, R. (2026). Human-AI teaming under fire: Lessons from Ukraine's human-in-the-loop combat AI systems. Journal of Strategic Security, 19(2), 67–100. https://doi.org/10.5038/1944-0472.19.2.2583

Vanderborght, R., & Nadibaidze, A. (2025). Military demonstrations as digital spectacles: How virtual presentations of AI decision-support systems shape perceptions of war and security. European Journal of International Security, 11(3), 399–418. https://doi.org/10.1017/eis.2025.10015

Wolf, R. (2025). Applying precautions in target verification with AI decision support systems. Israel Law Review, 58(2–3), 282–305. https://doi.org/10.1017/s0021223725100071

© IIARJournal. All rights reserved.