

INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS



Open Access, Refereed Journal Multi-Disciplinary
Peer Reviewed

www.ijlra.com

DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

**Copyright © International Journal for Legal Research & Analysis.
All rights reserved.**

ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

ALGORITHMIC ACCOUNTABILITY GAP: PROSECUTING UNPREDICTABILITY IN AUTONOMOUS WEAPONS SYSTEMS UNDER INTERNATIONAL CRIMINAL LAW.

AUTHORED BY - BHARADWAJA REDDY CHINTAKUNTA

ABSTRACT

This paper examines the incompatibility between machine learning-enabled autonomous weapons systems (AWS) and the Mens rea requirements of the Rome Statute, specifically Article 30's mandate for intent and knowledge. While international humanitarian law permits human-supervised autonomous weapons, the introduction of machine learning creates unpredictability that transcends human foreseeability, rendering traditional criminal responsibility frameworks useless. This research identifies the accountability vacuum and proposes innovations to preserve criminal liability in the era of algorithmic warfare.

The paper reveals a diffusion of culpability across the chain: (1) developers claim algorithmic unpredictability exempts them from liability; (2) operators argue lack of real-time control over autonomous decisions, and (3) commanders invoke reasonable reliance on certified systems. The Rome Statute's strict interpretation of Mens Rea, excluding recklessness and dolus eventualis doesn't cover prosecution when AWS commit indiscriminate attacks through machine learning behaviours rather than deliberate human intent.

Furthermore, Article 28's command responsibility doctrine fails to treat AWS as "subordinates," leaving violations without legally liable means. International humanitarian law's "meaningful human control" standard remains vague and cannot substitute for criminal law doctrine.

The paper advances a practical reform, it argues for the intervention of Appeals chamber or judicial innovation to expand the scope of Article 30 of the Rome Statute, on the ground that its strict intent and knowledge standard creates a danger of impunity when lethal decision-making is delegated to autonomous systems.

A. Introduction: Machine Learning and the Rome Statute

The increased deployment of machine learning-enabled Autonomous Weapons Systems (AWS) presents a structural challenge to the foundations of international criminal law. Emerging battlefields are being increasingly shaped by probabilistic decision-making, data-driven target recognition, and reinforcement learning processes that operate at machine speed and with opaque internal logic.¹ Yet the Rome Statute of the International Criminal Court still only presupposes human-controlled violence, committed by actors (humans) whose mental states can be reconstructed through the concepts of intent and knowledge.² The resulting friction between stochastic autonomy and anthropocentric culpability doctrines generates a phenomena called as the **Algorithmic Accountability Gap**.

How can the Rome Statute's current structure of mens rea and command responsibility be reinterpreted or doctrinally supplemented to preserve criminal accountability for any harm caused by machine learning-enabled autonomous weapons systems while remaining inclusive of the principle of legality and individual responsibility?

B. The Architecture of International Criminal Justice

I. International Criminal Law's Anthropocentric Design

International Criminal Law was constructed as a response to human atrocities orchestrated by leaders. The Nuremberg Trials established the principle that individuals at the apex of hierarchical command structures bore criminal responsibility for mass violence perpetrated by subordinates.³ This model was refined through the International Criminal Tribunals for the Former Yugoslavia (ICTY) and Rwanda (ICTR), which developed nuanced doctrines of command responsibility, co-perpetration, and complicity designed to fix liability on individuals whose decisions set criminal enterprises in motion.⁴ The Rome Statute, adopted in 1998, codified these doctrines in Articles 25 through 28, establishing ways of individual responsibility on three core assumptions: first, that violence is planned by human actors at a command centre; second, that violence is implemented by human subordinates operating

¹ Human Rights Watch, *A Hazard to Human Rights: Autonomous Weapons Systems and Digital Decision-Making* (2025) <https://www.hrw.org/report/2025/04/28/a-hazard-to-human-rights/autonomous-weapons-systems-and-digital-decision-making> accessed 19 December 2025.

² Rome Statute of the International Criminal Court, opened for signature 17 July 1998, 2187 UNTS 90 (entered into force 1 July 2002).

³ *Trial of the Major War Criminals before the International Military Tribunal*, Vol. XXII, p. 466 (Nuremberg 1948).

⁴ *Prosecutor v Tadić* (Decision on the Defence Motion for Interlocutory Appeal on Jurisdiction) IT-94-1-AR72 (ICTY Appeals Chamber, 2 October 1995).

within clear established hierarchies; and third, that responsibility can be concentrated at the apex of authority through doctrines of superior responsibility and co-perpetration.⁵

It was designed for a world in which commanders issued orders to human soldiers, who then executed those orders with greater or lesser fidelity. It was not designed for a world of semi-autonomous algorithms whose behaviour cannot be fully predicted.

II. Contemporary Conflict Trends

The contemporary landscape has shifted dramatically. Data from the Uppsala Conflict Data Program reveals that 2024 marked a historic high in the number of armed conflicts, with sixty-one active conflicts involving at least one state, the highest number since statistics began in 1946.⁶ If the ICC cannot adapt to address the primary means by which future wars will be fought, the consequence would be not merely institutional failure but a lack of accountability: low-technology actors would face scrutiny and prosecution, while high-technology powers would operate with impunity, creating a perverse incentive structure that wrongly rewards technological sophistication in war crimes.

C. From Automation to Autonomy

I. Distinguishing Deterministic Automation from Probabilistic Autonomy

To understand the legal challenge posed by AWS, one must first understand the differences between automation and autonomy in military systems. Traditional automated systems such as landmines or heat-seeking missiles, function on deterministic logic: if a specific triggering event occurs (pressure, thermal signature), a predetermined consequence follows (detonation, missile guidance). These systems exhibit mechanical or algorithmic predictability: their behaviour is rule-based, encoded in physical mechanisms or rigid conditional logic (“If X, then Y”). The outcome is foreseeable. If a landmine kills a civilian, it is because the landmine was placed in that location (attributing intent to the actor who placed it) or because the area was not properly warned or demarcated (negligence). The causal chain is linear; the responsibility locus is also identifiable.⁷

⁵ Rome Statute (n 2) arts 25-28.

⁶ Uppsala University, ‘UCDP: Sharp Increase in Conflicts and Wars’ (11 June 2025) <https://www.pcr.uu.se/research/ucdp/> accessed 14 December 2025.
Peace Research Institute Oslo (PRIO), ‘New Data Shows Conflict at Historic High’ (11 June 2025) <https://www.prio.org> accessed 15 December 2025.

⁷ International Committee of the Red Cross, *Autonomy, Meaningful Human Control and Weapons in International Humanitarian Law: Technical Perspectives* (2020).

Whereas, autonomous systems using machine learning operate on different principles. Rather than following pre-programmed rules, these systems are “trained” on vast datasets to optimize for a defined goal. Through processes of deep learning and reinforcement learning, the system develops internal representations, neural network weights and feature hierarchies that allow it to respond to situations without explicit programming of each contingency. The system exhibits behaviour that arise from the learned patterns but cannot be easily mapped to human made rules.⁸

The specific output from autonomous systems cannot be clearly foreseen by any individual developer, nor can it be outlined as a direct consequence of human programming. However, characterizing this on the face as ‘unpredictable’ obscures the nature of the risk and the precise locus of culpability. To properly assign mens rea, we must distinguish between **aleatoric uncertainty** and **epistemic uncertainty**.

Aleatoric uncertainty⁹ is the irreducible stochastic noise inherent in the battlefield environment that creates randomness in sensory data. Errors arising from this source are statistically foreseeable but individually unpredictable; deploying a system with known aleatoric risks implicates *dolus eventualis* or recklessness and the degree of aleatoric uncertainty can determine whether liability exists and to which extent?

Epistemic uncertainty¹⁰, by contrast, refers to reducible errors resulting from a lack of knowledge, such as model inadequacy, ‘bad data,’ or ‘out-of-distribution’ (OOD) deployment where the training data fails to match the operational environment. Errors arising from epistemic uncertainty are not ‘accidents’, they are failures of design and verification.¹¹ They implicate negligence and the ‘should have known’ standard of command responsibility, as the uncertainty was theoretically reducible through rigorous testing prior to deployment. This distinction pierces the ‘Black Box’ defence: if the error was epistemic, the commander failed in their duty of due diligence; if aleatoric, they consciously accepted a statistical certainty of harm.

⁸ Dongjo Kim, Jeongsu Lee and Ho-Young Kim, ‘Commanding Emergent Behavior with Neural Networks’ (2025) 6 Cell Reports Physical Science 102857.

⁹ Armen Der Kiureghian and Ove Ditlevsen, ‘Aleatory or Epistemic? Does it Matter?’ (2009) 31(2) Structural Safety 105.

¹⁰ Eyke Hüllermeier and Willem Waegeman, ‘Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods’ (2021) 110 Machine Learning 457.

¹¹ Ibid.

II. The Accountability Vacuum

States are rushing to develop AWS for compelling strategic reasons. They can cycle through the “OODA Loop” (Observe, Orient, Decide, Act) in milliseconds, far faster than human cognition.¹²

However, this speed and ease come at the cost of control. When a commander deploys a swarm of autonomous drones, the commander is not firing a gun; the commander is releasing a semi-independent agent. If that agent commits a war crime, attacks a hospital, misidentifies civilians as targets, causes disproportionate harm, the causal distance between the commander’s decision to “deploy” and the machine’s decision to “kill” creates a legal gap. Operators report a **“lack of real-time control”** over autonomous decisions. This phrase captures a real phenomenological and operational reality: as systems operate at machine speed, as their decision logic becomes opaque, and as the volume of targets and decisions is high it overwhelms human capacity to monitor each one, the operator’s role becomes a mere spectator.¹³

Psychological research confirms that humans interacting with automated systems exhibit a **“automation bias,”**¹⁴ which is a tendency to trust computer-generated recommendations over independent human judgment, especially when the computer system has previously been certified or validated.¹⁵ This creates a perfect **accountability vacuum**. Military commanders can argue that they deployed the system in compliance with legal certifications; the system operated within its parameters; the specific error was as a result of the learned model that no reasonable foresight could have anticipated.” Operators can argue: “I lacked real-time control; I trusted the certified system; the decision occurred too fast for me to intervene.” Developers can argue: “The system was designed to comply with IHL; the specific error resulted from environmental factors beyond our design parameters; the black-box nature of deep learning means that specific errors cannot be predicted.”¹⁶ Each actor points to another as the bearer of

¹² Stockholm International Peace Research Institute (SIPRI), ‘Compliance with International Humanitarian Law in the Development of Autonomous Weapons Systems’ (2023) <https://www.sipri.org> accessed 22 December 2025.

¹³ Antonio Coco, ‘Exploring the Impact of Automation Bias and Complacency on Individual Criminal Responsibility for War Crimes’ (2023) 21 JICJ 1077.

¹⁴ Mary L Cummings, ‘Automation Bias in Intelligent Time Critical Decision Support Systems’ (AIAA 1st Intelligent Systems Technical Conference, Chicago, September 2004).

¹⁵ Coco (n 13).

¹⁶ United Nations Institute for Disarmament Research (UNIDIR), ‘The Black Box, Unlocked: Methods for Determining Accountability in Autonomous Weapon Systems’ (2020) <https://unidir.org/publication/black-box-unlocked> accessed 20 December 2025.

responsibility, and the result is systematic non-accountability.

D. Doctrinal Crisis under the Rome Statute: Articles 30 and 28

I. Article 30

Article 30 of the Rome Statute serves as the default rule for the mental element (*mens rea*) required for most crimes within the ICC's jurisdiction. The provision states: "Unless otherwise provided, a person shall be criminally responsible and liable for punishment for a crime within the jurisdiction of the Court only if the material elements are committed with intent and knowledge."¹⁷ This represents a **strict mens rea standard**, excluding recklessness and *dolus eventualis* as standalone bases of liability.¹⁸

Article 30(2) defines intent in two alternative formulations. First: "In relation to conduct, the person means to engage in the conduct." This captures direct intent to act. Second, and more critically for AWS scenarios: "In relation to a consequence, the person means to cause that consequence **or is aware that it will occur in the ordinary course of events.**"¹⁹ This second formulation reflects the civil law concept of indirect intent: a person intends a consequence not because they desired it as a primary goal, but because they foresaw it as a virtually certain outcome of their conduct and proceeded anyway, accepting that consequence as a collateral result.

The phrase "**in the ordinary course of events**" is contentious. In the case *Prosecutor v. Lubanga* (2007)²⁰, the Pre-Trial Chamber offered an interpretation incorporating elements of *dolus eventualis* (recklessness or acceptance of risk): the Court suggested that intent could be satisfied where the accused "is aware of the risk that the objective elements of the crime may result from his or her actions or omissions, and (b) accepts such an outcome by reconciling himself or herself with it or consenting to it." However, this interpretation was substantially narrowed in *Prosecutor v. Bemba* (2009)²¹, where the Pre-Trial Chamber emphasized that the threshold for "ordinary course of events" requires "**virtual certainty**" or "**practical certainty**" of the consequence, not merely substantial likelihood or accepted risk. The 2014

¹⁷ Rome Statute (n 2) art 30.

¹⁸ Sara Porro, 'ICC Statute Article 30(2)(b) Interpretation' (2020) <https://www.legal-tools.org> accessed 12 December 2025.

¹⁹ Ibid.

²⁰ *Prosecutor v. Lubanga Dyilo* (Pre-Trial Chamber Decision) ICC-01/04-01/06 (29 January 2007).

²¹ *Prosecutor v Jean-Pierre Bemba Gombo* (2009) ICC-01/05-01/08-424 [362]–[363].

Lubanga Appeals Judgment²² confirmed this restrictive standard: “the standard for the foreseeability of events is virtual certainty.”

To understand the gap, consider a hypothetical but realistic scenario involving a military commander who deploys an Autonomous Weapon System (AWS) to patrol a contested border region. In this scenario, the commander fully intends to deploy the system, thereby satisfying the conduct element of the crime. However, the system malfunctions or exhibits a learned error, misidentifying a refugee group sheltering in the area as a militia unit and engaging them with lethal force.

Prosecutors seeking to hold the commander accountable now face a evidentiary burden: they must prove that the commander, at the precise moment of deployment, was aware with “virtual certainty” that this specific civilians in this specific location, under these specific environmental conditions would be targeted and harmed. This standard is nearly impossible to meet due to the statistical nature of AWS performance.²³ The appeal of AWS to military planners is precisely that they achieve high success rates while accepting a calculable with a small, error rate.²⁴

This dynamic creates a perverse legal incentive. If an AWS system is sufficiently accurate, for instance, a 95 to 99 percent success rate, the rare errors fall below the threshold of what occurs in the “ordinary course of events.” Therefore, no commander can be said to have foreseen these specific errors with virtual certainty. Yet, these rare errors are not negligible; when they occur in the context of widespread deployment across a war, they will accumulate to cause significant, although statistically “rare,” civilian harm.

Article 30(3) of the Rome Statute defines knowledge as “awareness that a circumstance exists or a consequence will occur in the ordinary course of events.”²⁵ Superficially, this provision might be read to include algorithmic harms: if a commander deploys a system with known error rates in a civilian-dense area, one could argue the commander “knows” that consequences (civilian harm) will occur. However, the “black box”²⁶ problem inherent in machine learning

²² *Prosecutor v. Lubanga Dyilo* (Appeals Chamber Judgment) ICC-01/04-01/06 A (1 December 2014).

²³ Human Rights Watch (n 1).

²⁴ *Ibid.*

²⁵ Rome Statute (n 2) art 30(3).

²⁶ UNIDIR (n 16).

hinders this understanding significantly.

Deep neural networks are often deeply opaque; even their creators cannot explain the specific decision-making process that led to a particular outcome.²⁷

The lack of knowledge about how the system works, creates a legal gap: the lack of liability for specific outputs.²⁸ Under current doctrine, the commander's "knowledge" is strictly limited to what is knowable through these standard procedures. Because the opaque nature of the learned model means that specific error modes are, by definition, not knowable ex ante, the commander can plausibly claim: "I had no knowledge of this specific error; the system was certified; the error resulted from environmental factors beyond my knowledge or control."

II. Article 28: Command Responsibility

When direct intent under Article 30 fails, prosecutors resort to invoke Command Responsibility, codified in Article 28 of the Rome Statute. This doctrine holds superiors liable for crimes committed by forces under their effective command and control if the superior knew or should have known of the crimes and failed to prevent or punish them.²⁹ Designed to also extend liability upward through hierarchical structures, this doctrine ensures that those at the apex of authority cannot escape responsibility by claiming ignorance. However, the doctrine was designed for human military hierarchies, not for human-machine hybrids.³⁰

Can an algorithm legally be treated as a "subordinate"? The answer appears to be no.³¹ It is not a subordinate in the legal sense because it lacks personhood, the capacity for intent, the capacity for obedience or disobedience, and moral agency. It cannot be ordered, cannot refuse an order, and cannot be held independently responsible.³²

²⁷ Samuel Stevens and others, 'Interpretable Algorithmic Forensics' (2023) 120 Proceedings of the National Academy of Sciences e2304245120.

²⁸ Human Rights Watch (n 1).

²⁹ Mark Klamberg, 'Article 28: Responsibility of commanders and other superiors' in Mark Klamberg (ed), Commentary on the Law of the International Criminal Court (Case Matrix Network 2017) <https://www.casematrixnetwork.org/cmn-knowledge-hub/clicc/> accessed 17 December 2025.

³⁰ Human Rights Watch (n 1).

³¹ Morten Bergsmo and Emilie Hunter, *Command Responsibility: International Criminal Law Guidelines* (Case Matrix Network, 2016).

³² Ibid.

As scholars like Jack Beard³³ and Tim McFarland³⁴ note, the superior-subordinate relationship in ICL is decided not just on causal control, but on the capacity to ‘prevent or punish.’ One cannot ‘punish’ an algorithm in a manner that satisfies the retributive and deterrent aims of the Rome Statute. Furthermore, the concept of ‘effective control’ relies on the moral authority and disciplinary mechanisms inherent to human hierarchy and the duty of obedience which are absent in human-machine interaction. Treating an AWS as a subordinate is a ‘functionalist trap’ that allows commanders to externalize responsibility to an entity that cannot be held liable, thereby deepening the accountability gap. Instead, the AWS must be legally conceptualized as a weapon or tool, shifting the focus to the commander’s direct individual liability for its reckless or negligent employment.

E. The Way Forward.

I. Article 121 and the Trap of Fragmentation

Proposals to amend the Rome Statute to specifically criminalize the reckless use of AWS or to create a new mode of liability for algorithmic harms are the first to come to mind. However, after analysis of the Rome Statute’s amendment provisions, combined with a realist assessment of the current geopolitical landscape, reveals that this path is practically useless.

“Article 121³⁵ of the Rome Statute governs the amendment process, establishing a high threshold for modifying the treaty that was deliberately designed to ensure stability and protect state sovereignty. The procedural hurdles embedded in this article are immense. An amendment requires adoption by a two-thirds majority of the Assembly of States Parties (ASP) or by a Review Conference. Under Article 121(4)³⁶, an amendment enters into force for all States Parties only one year after seven-eighths of them have deposited their instruments of ratification or acceptance.”

This “seven-eighths” requirement acts as a virtual lock on the Statute’s general provisions. With over 120 States Parties, achieving near-unanimous ratification involves an ideal degree of coordinating. The history of international treaty law implies that such a threshold is a

³³ Jack M Beard, ‘Autonomous Weapons and Human Responsibilities’ (2014) 45 *Georgetown Journal of International Law* 617 <https://digitalcommons.unl.edu/lawfacpub/196/> accessed 19 December 2025

³⁴ Tim McFarland and Tim McCormack, ‘Mind the Gap: Can Developers of Autonomous Weapons Systems be Liable for War Crimes?’ (2014) 90 *International Law Studies* 361 <https://digitalcommons.usnwc.edu/ils/vol90/iss1/2/> accessed 10 December 2025

³⁵ Rome Statute (n 2) art 121.

³⁶ *Ibid* art 121(4).

generational endeavour. In the context of AWS, By the time an amendment achieves the necessary ratifications, the technology it seeks to regulate will likely have evolved beyond the definitions contained in the text.

More critical than the ratification threshold is the specific regime governing amendments to crimes, set out in Article 121(5)³⁷. It states that any amendment to Articles 5, 6, 7, and 8³⁸ (the substantive crimes) enters into force only for those States Parties that have accepted the amendment. Crucially, the Court “shall not exercise its jurisdiction regarding a crime covered by the amendment when committed by nationals of or on the territory of a State Party which has not accepted the amendment.”

This mechanism creates a “fragmented jurisdiction”, it would essentially function as an opt-in protocol. States that are most active in developing and deploying AWS, primarily major military powers would simply refuse to ratify the amendment.

By creating a specific crime for AWS misconduct that major powers opt out of, the amendment would shield them from prosecution under the existing general provisions of the Statute.

II. Judicial Interpretation: The “Macro-Statistical” Approach

If the ICC is to address the challenge of algorithmic warfare, it must do so through the tools it already possesses: the judicial interpretation of the existing Statute. Current ICL doctrine struggles with both aleatoric and epistemic uncertainty. The *Bemba*³⁹ standard treats aleatoric uncertainty (e.g., a 5% error rate) as a “mere possibility,” insufficient for liability, and epistemic uncertainty (the black box problem) as a defence. By strictly requiring intent/knowledge, Article 30⁴⁰ allows commanders to hide behind both.

At the micro-level, a single decision by an AWS to fire has a low probability of error (e.g., 0.01%). A commander looking at that single engagement cannot be said to have “virtual certainty” that a crime will occur. However, at the macro-level, the laws of probability, specifically the Law of Large Numbers⁴¹ dictate that as the number of engagements increases,

³⁷ Ibid art 121(5).

³⁸ Ibid art 5,6,7,8.

³⁹ *Bemba* (n 21)

⁴⁰ Rome Statute (n 2) art 30.

⁴¹ William Feller, *An Introduction to Probability Theory and Its Applications* (3rd edn, Wiley 1968) 228.

the cumulative probability of error approaches 100%. If an AWS with a 99% accuracy rate engages 1,000 targets, it is a statistical virtual certainty that it will commit roughly 10 errors.

While a single autonomous decision carries only a low probability of error, the aggregated deployment of a stochastic system over thousands of engagements raises that probability to a **Statistical Virtual Certainty**. The “ordinary course of events” for a stochastic system is not a single straight line of cause and effect, but a probability distribution. When that distribution dictates that a certain volume of civilian harm is mathematically inevitable given the scale of deployment, the “virtual certainty” threshold of *Bemba* is met at the aggregate level.

Legal Logic of the Macro-Statistical Approach

1. **Deployment:** Commander X deploys a swarm of AWS.
2. **Error Profile:** The system has a known error profile (e.g., 1% false positive rate).
3. **Scale:** The scale of deployment (1,000 engagements) makes it a statistical virtual certainty (approx. 100%) that civilian casualties will occur.
4. **Culpability:** Therefore, Commander X is “aware that civilian casualties will occur in the ordinary course of events” (Article 30(2)(b)).

The fact that the commander does not know *which* specific civilians will be killed is irrelevant; they know that the consequence (death of protected persons) is certain. This interpretation aligns with the *Bemba*⁴² requirement for “certainty” while stripping away the defence that the specific instance was uncertain. It shifts the *mens rea* object from the “specific victim” to the “class of victims” harmed by the system’s operation. It acknowledges that in algorithmic warfare, certainty is a property of the aggregate, not the individual event.

III. Risk-Based Delegation

The second pillar of this reform is to recognize the act of delegation as the locus of culpability. Current ICL focuses on the moment of the “attack.” However, with AWS, the human involvement ends at the moment of deployment. The Court should recognize that delegating lethal authority to a system with known aleatoric risks in an environment of civilian-combatant ambiguity constitutes a culpable mental act. This is not “negligence” (which is excluded by Art 30⁴³); it is a form of intent constructed through the conscious acceptance of statistical inevitability.

⁴² *Bemba* (n 21)

⁴³ Rome Statute (n 2) art 30.

To respect the principle of legality and prevent this from becoming a strict liability standard, the Court should establish specific criteria for when delegation becomes culpable:

- **Operational Context:** The environment exhibits high ambiguity (e.g., dense urban terrain vs. open desert) where the system's error rate is likely to manifest as civilian harm.
- **Known Failure Modes:** The system has known performance deficits relevant to that context (e.g., data showing the algorithm struggles to distinguish non-combatants in specific lighting conditions).
- **Intelligence Possession:** The commander possessed intelligence regarding these risks and the environmental mismatch yet proceeded with delegation.

Under this standard, the “strategic choice” to externalize risk is to sacrifice precision for the speed and scale of automation becomes the criminal act. By focusing on the error rate rather than the error mechanism, the Court avoids getting bogged down in technical debates about explainable AI (XAI) and focuses on the commander's decision to unleash a system with a known statistical propensity for harm. If a commander releases a “Black Box” that kills 10% of the time, they cannot claim lack of intent simply because they don't understand the internal wiring. The output, the 10% death rate is part of the “ordinary course of events” for that system.

IV. Harmonious co-existence with the Rome Statute

Critics may argue that this approach lowers the *mens rea* threshold, smuggling recklessness (*dolus eventualis*) back into Article 30 contrary to the drafter's intent. However, this reform is strictly consistent with the teleological purpose of the Statute and the rules of interpretation mandated by international law.

Specifically, this interpretation accords with **Article 31(1) of the Vienna Convention on the Law of Treaties (VCLT)**⁴⁴, which dictates that a treaty shall be interpreted “in good faith in accordance with the ordinary meaning to be given to the terms of the treaty in their context and in the light of its object and purpose.” The object and purpose of the Rome Statute, affirmed in its Preamble, is to put an end to impunity for the most serious crimes of concern to the international community. A rigid, static interpretation of “ordinary course of events” that ignores the mathematical realities of algorithmic warfare would effectively grant immunity to high-technology powers, thereby defeating the Statute's very object and purpose. Therefore,

⁴⁴ *Vienna Convention on the Law of Treaties* (adopted 23 May 1969, entered into force 27 January 1980) 1155 UNTS 331 (VCLT) art 31(1).

adopting the “Macro-Statistical” approach is not an expansion of the law, but a necessary evolutionary interpretation required by the VCLT to prevent the Statute from becoming obsolete in the face of technological change.

The interpretation of “ordinary course of events” remains a matter of judicial construction, not statutory rewriting. The *Bemba* standard itself was a judicial innovation that narrowed the text; broadening it to accommodate statistical reality is equally within the Chamber’s mandate. Furthermore, “Statistical Virtual Certainty” is not probabilistic recklessness in the traditional sense. Recklessness involves taking a “substantial risk”. In contrast, Statistical Virtual Certainty involves proceeding where the accumulation of risk makes the outcome inevitable. The Macro-Statistical approach respects the Article 30 exclusion of *dolus eventualis* for single acts but recognizes that aggregated acts transform risk into knowledge.

F. Conclusion

The proliferation of machine learning-enabled Autonomous Weapons Systems (AWS) threatens to make the Rome Statute’s frameworks obsolete, creating an accountability vacuum. As demonstrated, existing doctrines under Articles 30 and 28 fail to adequately address the specific nature of algorithmic unpredictability and the black box problem. Furthermore, reliance on legislative amendments via Article 121 is politically unviable and risks jurisdictional fragmentation. Therefore, preserving criminal liability requires a judicial pivot toward a Macro-Statistical interpretation of *mens rea*. By reframing the “ordinary course of events” to recognize “**Statistical Virtual Certainty**” in aggregated deployments, the International Criminal Court can bridge the gap between stochastic autonomy and legal accountability. This approach ensures that the strategic delegation of lethal authority to risky algorithms remains a culpable act, preventing high-tech warfare from operating with impunity.

Bibliography

Table of Cases

- “*Prosecutor v Jean-Pierre Bemba Gombo* (Decision on the Confirmation of Charges) ICC-01/05-01/08-424 (15 June 2009)”
- “*Prosecutor v Lubanga Dyilo* (Decision on the Confirmation of Charges) ICC-01/04-01/06 (29 January 2007)”

- “*Prosecutor v Lubanga Dyilo* (Judgment on the Appeal against the Conviction) ICC-01/04-01/06 A (1 December 2014)”
- “*Prosecutor v Tadić* (Decision on the Defence Motion for Interlocutory Appeal on Jurisdiction) IT-94-1-AR72 (2 October 1995)”
- “*Trial of the Major War Criminals before the International Military Tribunal*, Vol XXII (Nuremberg 1948)”

Table of Legislation

- “Rome Statute of the International Criminal Court (opened for signature 17 July 1998, entered into force 1 July 2002) 2187 UNTS 90”
- “Vienna Convention on the Law of Treaties (adopted 23 May 1969, entered into force 27 January 1980) 1155 UNTS 331”

List of References

- Beard, Jack M, ‘Autonomous Weapons and Human Responsibilities’ (2014) 45 *Georgetown Journal of International Law* 617 <https://digitalcommons.unl.edu/lawfacpub/196/> accessed 19 December 2025
- Bergsmo, Morten and Hunter, Emilie, *Command Responsibility: International Criminal Law Guidelines* (Case Matrix Network 2016)
- Coco, Antonio, ‘Exploring the Impact of Automation Bias and Complacency on Individual Criminal Responsibility for War Crimes’ (2023) 21 *Journal of International Criminal Justice* 1077
- Cummings, Mary L, ‘Automation Bias in Intelligent Time Critical Decision Support Systems’ (AIAA 1st Intelligent Systems Technical Conference, Chicago, September 2004)
- Der Kiureghian, Armen and Ditlevsen, Ove, ‘Aleatory or Epistemic? Does it Matter?’ (2009) 31 *Structural Safety* 105
- Feller, William, *An Introduction to Probability Theory and Its Applications* (3rd edn, Wiley 1968)
- Human Rights Watch, *A Hazard to Human Rights: Autonomous Weapons Systems and Digital Decision-Making* (2025) <https://www.hrw.org/report/2025/04/28/a-hazard-to-human-rights/autonomous-weapons-systems-and-digital-decision-making> accessed 19 December 2025

- Hüllermeier, Eyke and Waegeman, Willem, 'Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods' (2021) 110 *Machine Learning* 457
- International Committee of the Red Cross, *Autonomy, Meaningful Human Control and Weapons in International Humanitarian Law: Technical Perspectives* (2020)
- Kim, Dongjo, Lee, Jeongsu and Kim, Ho-Young, 'Commanding Emergent Behavior with Neural Networks' (2025) 6 *Cell Reports Physical Science* 102857
- Klamberg, Mark, 'Article 28: Responsibility of commanders and other superiors' in Mark Klamberg (ed), *Commentary on the Law of the International Criminal Court* (Case Matrix Network 2017) <https://www.casematrixnetwork.org/cmn-knowledge-hub/clicc/> accessed 17 December 2025
- McFarland, Tim and McCormack, Tim, 'Mind the Gap: Can Developers of Autonomous Weapons Systems be Liable for War Crimes?' (2014) 90 *International Law Studies* 361 <https://digital-commons.usnwc.edu/ils/vol90/iss1/2/> accessed 10 December 2025
- Peace Research Institute Oslo (PRIO), 'New Data Shows Conflict at Historic High' (11 June 2025) <https://www.prio.org> accessed 15 December 2025
- Porro, Sara, 'ICC Statute Article 30(2)(b) Interpretation' (2020) <https://www.legal-tools.org> accessed 12 December 2025
- Stevens, Samuel and others, 'Interpretable Algorithmic Forensics' (2023) 120 *Proceedings of the National Academy of Sciences* e2304245120
- Stockholm International Peace Research Institute (SIPRI), *Compliance with International Humanitarian Law in the Development of Autonomous Weapons Systems* (2023) <https://www.sipri.org> accessed 22 December 2025
- United Nations Institute for Disarmament Research (UNIDIR), *The Black Box, Unlocked: Methods for Determining Accountability in Autonomous Weapon Systems* (2020) <https://unidir.org/publication/black-box-unlocked> accessed 20 December 2025
- Uppsala University, 'UCDP: Sharp Increase in Conflicts and Wars' (11 June 2025) <https://www.pcr.uu.se/research/ucdp/> accessed 14 December 2025