Judgement Day

The Ethics of AI in Warfare and the Future of Human Command

Published By: Ambient Stratagem

May 2025

Author: John Blamire

War, once a dominion of men, maps, and steel, is now shaped by silicon and signal. In 2025, global security is undergoing a profound transformation, not just in the weapons being developed, but in the logic that governs their use. Artificial intelligence is no longer an experimental adjunct to military power—it is fast becoming its core nervous system.

Yet amid this accelerating shift, one element remains constant: the ethical responsibility for life and death in conflict.

The Strategic Context – The Battlefield Has Changed, but Ethics Have Not

The Erosion of International Norms

The rules-based order that emerged after 1945—anchored in the Geneva Conventions, international humanitarian law, and a shared moral vocabulary among democracies—is being unpicked. Russia's war on Ukraine, Chinese cyber-aggression, and Iran's proxy network tactics have all demonstrated how peer adversaries exploit legal and ethical ambiguity. AI-enabled warfare threatens to deepen this erosion by creating "plausible deniability by design"—striking without clear attribution or intent.

Fragmentation of Alliance Trust

NATO faces internal frictions: budget shortfalls, asymmetric threat perception between Western and Eastern members, and rising concern over US reliability. Washington's shift toward Indo-Pacific priorities, and domestic isolationist currents, have shaken confidence among European partners. The result is an emerging doctrine of differentiated ethics—where national AI deployment strategies may diverge even within allied structures.

Technological Acceleration Outpacing Doctrine

Loitering munitions with AI-driven target recognition, semi-autonomous unmanned ground vehicles, and decision-support systems that operate in milliseconds now populate battlefields. Yet formal doctrine remains slow to adapt. The speed of AI-facilitated engagements challenges not only legal review processes but the very notion of command accountability.

> "The great trap of our age is believing speed is virtue. But in war, unchecked speed often precedes moral collapse." — General Sir Rupert Smith, The Utility of Force

AI and the Moral Compression of War

The appeal of AI in warfare is seductive: faster decisions, greater precision, fewer humans at risk. But these perceived gains come with a hidden cost—the compression of time for ethical reflection. Human judgement, historically central to the laws of war, risks becoming an afterthought in machine-accelerated battlespaces.

This is not hypothetical. In the 2023 Nagorno-Karabakh flashpoint, Turkish-supplied drones operated with AI-enhanced pattern recognition, reportedly targeting movement without live operator intervention. While hailed as a tactical success, reports of civilian casualties—misidentified due to ambiguous thermal signatures—have raised profound questions over accountability.

And it's not just in the Global South. In Ukraine, both NATO and Russian-aligned forces have deployed semi-autonomous ISR systems with loiter-and-strike capability. The battlefield is becoming a proving ground for autonomy—at the expense of clarity.

The Urgency of Ethical Reaffirmation

There is a critical window—measured in years, not decades—in which NATO and its partners can set the moral terms for AI warfare. To do so, they must:

- Reassert the primacy of human oversight, even in edge environments where comms are denied or degraded.
- Refuse the logic of "technological determinism"—the idea that because something can be automated, it should be.
- Build interoperability not just in hardware, but in values—ensuring ethical congruence across allied systems.

Because once ethical leadership is lost, it cannot be regained in wartime.

> "In the absence of shared moral boundaries, war becomes a test not of right, but of reach." — Lt. Gen. H.R. McMaster (Ret'd)

What Is Ethical AI in Warfare?

Accountability

Every decision—especially lethal ones—must be traceable to a human authority

Auditability

Decisions made by AI systems must be reconstructable and reviewable

Explainability

Operational users must understand the AI's recommendation in mission-relevant timescale

Governability

Human operators must retain ultimate control

As machines edge closer to making life-and-death decisions on the battlefield, the question is no longer whether we can build autonomous combat systems—but whether we should. And if we do, how do we ensure those systems act within the boundaries of military ethics, international law, and human conscience?

The concept of ethical AI in warfare is not merely philosophical—it is operational, legal, and deeply strategic. It determines how military decisions are made, who is accountable, and whether legitimacy is preserved when machines act faster than humans can think.

> "The more complex the machine, the more necessary the human conscience becomes." — General Sir David Richards, Former Chief of the Defence Staff

NATO's Baseline: Meaningful Human Control

NATO's Principle

In 2021, NATO adopted a principle that AI systems must operate under "meaningful human control." But in 2025, this phrase is under strain. What is meaningful in a denied environment? Can a drone swarm controlled by predictive algorithms still be meaningfully governed by a single human on a battlefield?

Without further doctrinal clarity, this principle risks becoming rhetorical cover rather than ethical safeguard.

The Edge Dilemma: Speed vs Morality

In modern warfare, especially in contested electronic environments, human–machine interaction is often reduced to milliseconds. A loitering munition may identify a heat signature, correlate it to a known pattern, and initiate a strike before any human operator can intervene.

This creates a fundamental dilemma:

- Delay the strike, and risk mission failure or fratricide.
- Authorise autonomous action, and risk disproportionate force, civilian harm, or war crimes.

Ethical Latency

AI systems operating at the edge must therefore be built not for speed alone—but for ethical latency. That is: systems should prioritise human-in-the-loop approval when ambiguity is high and operate autonomously only under tightly bounded conditions.

Ethical latency is not a brake—it is a guardrail.

> "In the end, the decision to kill must be one a human owns—morally, legally, and historically." — Mary Wareham, Human Rights Watch (on LAWS)

The Moral Dilemmas of Autonomous Weapons

The moment a machine is authorised to take human life without further human consultation marks a fundamental rupture in the ethics of warfare. It is not simply a technical milestone—it is a moral fracture. Autonomous weapons, by their very design, challenge the philosophical, legal, and emotional foundations of how we understand violence, responsibility, and justice in war.

This chapter confronts the dilemmas that emerge when algorithms assume the role of executioner.

The Distinction Problem

At the heart of international humanitarian law lies the principle of distinction—the obligation to differentiate between combatants and non-combatants. In urban, irregular, or asymmetric warfare, this is already a challenge for human soldiers.

For AI, it is exponentially harder.

- How does a drone interpret a child carrying a metal pipe?
- Can a surveillance algorithm distinguish between a farmer returning at night and an insurgent planting an IED?
- Will facial recognition systems misidentify ethnic minorities or the disabled based on flawed training data?

These are not theoretical. In Gaza (2023), and Eastern Ukraine (2024), autonomous systems were reportedly involved in strikes where civilian harm was later contested but attribution was inconclusive—raising questions that remain unanswered.

> "Machines have no understanding of context, no instinct to pause. That pause is often where humanity lives." — Professor Noel Sharkey, AI ethicist and former MoD advisor

The Problem of Machine Error

False Positives

Misidentifying neutral or friendly forces as hostile

Overfitting

AI trained in one combat environment misapplies logic in another

Sensor Fusion Failure

Conflicting data from thermal, acoustic, or visual feeds

Black Swan Events

Behaviour in novel environments that has no precedent in training data

All weapons systems fail. But when autonomous weapons fail, they fail differently—and more dangerously. A misfiring rifle affects a moment. A misclassifying algorithm can result in dozens dead.

And unlike human error, machine error is often non-obvious. It may take months or years before a flawed targeting algorithm is discovered—by which point, the consequences may be irreversible.

The Accountability Vacuum

The Developer

Who wrote the targeting algorithm?

The Data Scientist

Who labelled the training images?

The Commander

Who authorised the system's use?

The Defence Ministry

Who procured the system?

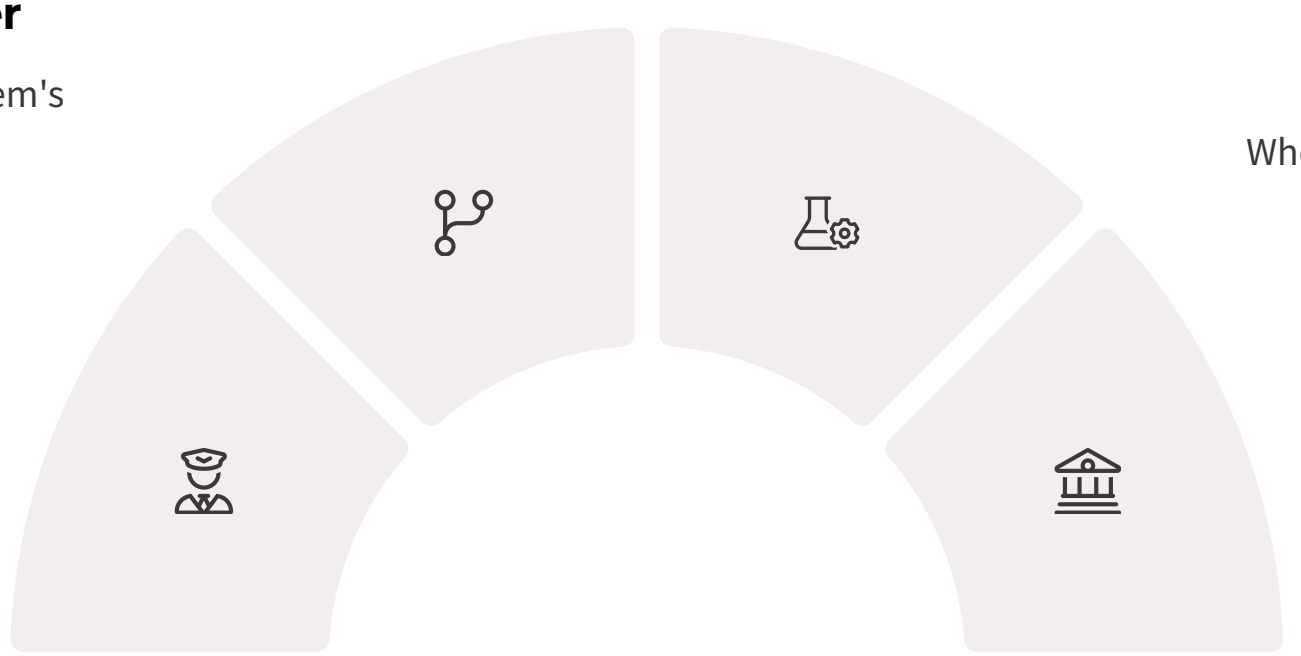

If an autonomous weapon kills unlawfully, who is accountable?

No military alliance—NATO included—has yet resolved this question in law or practice. The absence of a clear accountability chain for machine-led killing is not just an ethical gap; it is a strategic liability. It opens the door to enemy propaganda, legal retaliation, and erosion of allied legitimacy.

This is compounded by the reality that:

- Peer adversaries (notably Russia and China) operate with vastly different ethical frameworks
- Non-state actors are not bound by Geneva obligations at all
- Civil society groups are mobilising global campaigns for bans or moratoria

> "You cannot court-martial an algorithm. But you can lose the war of legitimacy because of one." — Ret. Gen. Richard D. Clarke, Former SOCOM Commander

The Escalation Risk

Tactical Overuse

Deploying autonomous systems in environments where humans would hesitate

Strategic Miscalculation

Triggering retaliation based on unintended autonomous escalation

Moral Disengagement

Diluting institutional responsibility within military chains of command

Autonomous weapons reduce the cognitive and emotional cost of killing. This detachment creates a risk of escalation through multiple pathways.

One senior NATO planner recently warned: "What we gain in tempo, we risk losing in restraint."

As conflict becomes increasingly algorithmic, the risk is not just that we lose control—but that we forget what it means to have had it in the first place.

NATO vs Adversaries – Diverging Ethical Doctrines

NATO	China	Russia
Governable Systems	Speed Over Scrutiny	Operational Autonomy Without Oversight
Explainability & Transparency	Centralised Control, Tactical Autonomy	Doctrine of Ambiguity
Legal Compliance	No Formal Ethical Doctrine	Disregard for Civilian Distinction
Ethical Consistency	Strategic Opportunism	Plausible Deniability
Alliance Interoperability	CCP Oversight	Strategic Fog

Military ethics are not universal. They are culturally grounded, politically reinforced, and strategically shaped. In the race to integrate artificial intelligence into warfare, NATO and its peer adversaries—especially China and Russia—are now on visibly divergent ethical trajectories.

NATO's official stance, reaffirmed in its updated 2024 AI Strategy, is grounded in the principle of "Responsible AI". This approach reflects NATO's deeper identity as a political alliance, not merely a military one. Ethics are a binding agent. To violate them for tactical advantage is to threaten the strategic cohesion of the alliance itself.

> "What separates us from our adversaries is not our firepower—it is our restraint." — Jens Stoltenberg, NATO Secretary General

China: Strategic Opportunism and Tactical Autonomy

Speed Over Scrutiny

China prioritises decision superiority. It views the capacity to strike faster than the adversary as ethically justified by the presumed deterrent effect.

Centralised Control, Tactical Autonomy

While strategic oversight is tightly held by the CCP, battlefield AI systems are being designed with increasing levels of real-time autonomy—particularly in drone swarms, missile targeting, and cyber-electronic warfare.

No Formal Ethical Doctrine

Unlike NATO, China has not publicly articulated ethical limits to autonomous weapons deployment. Its regulatory frameworks focus on effectiveness, not moral constraint.

China's People's Liberation Army (PLA) sees AI not as an ethical challenge, but as a strategic multiplier. Its doctrine—articulated through sources like the Science of Military Strategy and internal PLA briefings—highlights these key beliefs.

Chinese state media has openly promoted AI as a way to offset the "human weaknesses" of conventional forces, describing future conflicts as a contest of "machine minds."

Russia: Reckless Innovation, Moral Ambiguity

Operational Autonomy Without Oversight

Minimal human input in targeting decisions

Doctrine of Ambiguity

Embracing "plausible deniability" and "strategic fog"

Disregard for Civilian Distinction

Training AI with problematic targeting data

Russia's military-industrial AI effort has been accelerated by its war in Ukraine. With constrained manpower and sanctions limiting hardware imports, Moscow has turned to AI as a way to level the playing field.

Russian forces have tested AI-enhanced ISR drones and loitering munitions that identify targets with minimal human input. In some cases, strikes have occurred without verifiable human authorisation.

Russian information warfare doctrine embraces "plausible deniability" and "strategic fog." AI systems that obscure attribution or accountability fit comfortably within this worldview.

In Syria and Ukraine, Russia has repeatedly been accused of targeting civilian infrastructure. AI-enabled systems, critics argue, are being trained with such data—baking in a disregard for legal norms.

> "Russia's approach to AI in warfare is not governed by Geneva, but by Gerasimov." — Col. John Spencer, Modern War Institute

The Risk of Ethical Asymmetry

27%

Global Support

Only 27% of respondents supported autonomous weapons without human control

61%

Global Opposition

Percentage of global respondents who opposed autonomous weapons outright

3

Ethical Frameworks

Major divergent approaches to AI ethics between NATO, China, and Russia

The divergence between NATO and its adversaries creates a strategic paradox:

- NATO adheres to principles that may slow its tempo and restrict the use of certain AI capabilities.
- Adversaries, unbound by such constraints, may deploy more aggressive and risk-tolerant systems.

This leads to a scenario where:

- NATO may appear operationally cautious, even when morally correct.
- Adversaries may gain tactical advantage, even at the cost of legitimacy.
- Public trust within NATO nations may erode if adversaries exploit these asymmetries for propaganda.

The worst-case outcome is not just a battlefield loss—it is the erosion of the West's moral leadership in warfare.

Designing Ethical Human–AI Teaming

Cognitive Integration

The human must understand what the machine knows, and vice versa

Shared Intent

AI must align with commander's goals, not just optimal efficiency

Trust Calibration

Systems must be tested, reviewed, and explained continually

Ethical Fail-safes

Default to human review when judgement is ambiguous

If the battlefield of the future is to remain governed by moral law and strategic control, then its systems must be built not for pure autonomy—but for ethical collaboration. The answer is not to ban AI from warfare, nor to surrender command to machines, but to design systems where human and machine work together—seamlessly, intelligently, and accountably.

This is the heart of human–AI teaming.

Historically, military technologies—from the longbow to the jet engine—have acted as tools: amplifying human capability, but always under direct command. AI systems change this dynamic. They interpret, decide, and act—creating the potential for systems that co-author military outcomes.

> "If AI is a fellow soldier, then it must be governed by the same rules, the same values, and the same ability to be held to account." — Air Chief Marshal Sir Mike Wigston, RAF (Ret'd)

Design Principles for Ethical Teaming

Explainable Interfaces

Operators must be able to interrogate AI logic —"Why did you suggest this target?"—and receive answers in real-time, appropriate to the command level.

Layered Autonomy

Autonomy should vary by function, context, and theatre. For example: logistics optimisation may be highly autonomous, while lethal targeting must demand higher thresholds for human validation.

Distributed Oversight

Teams should include ethical oversight roles—either embedded within units or remote via secure comms —who monitor AI decisions for compliance with rules of engagement and international law.

Edge Survivability, Not Ethical Abdication

In degraded or denied environments, systems must retain ethical constraints locally. This includes pre-loaded rules of engagement, failsafe disengagement logic, and hierarchical escalation management.

Continuous Human Trust Calibration

Trust is not static. Systems must be tested, reviewed, and explained continually to maintain operator confidence. Trust without understanding is dangerous. But so is distrust without cause.

To embed ethics into human–AI teams, NATO systems must be designed with mission-first and value-bound architecture.

> "In a machine-accelerated world, ethical delay must be a feature—not a flaw." — Lt. Gen. James Hockenhull, Commander, UK Strategic Command

Public Consent and the Legal Fragility of Autonomous Warfare

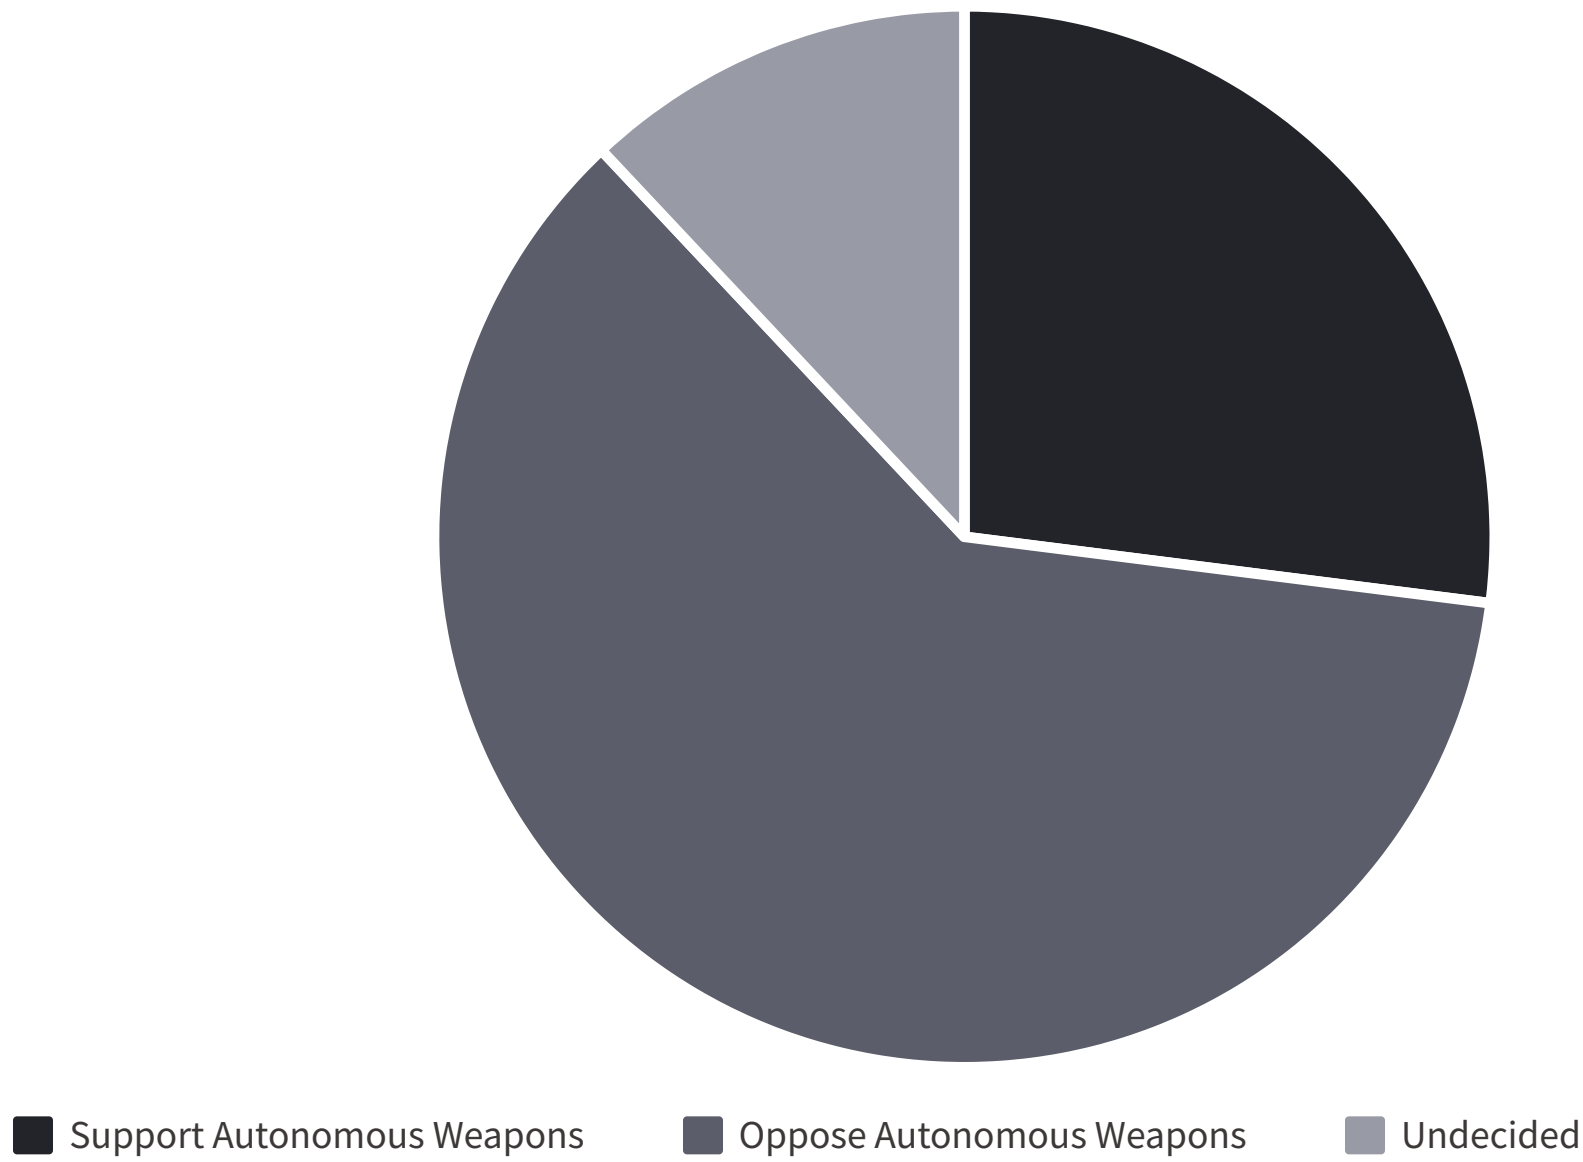

Military power in democracies rests not only on technological superiority, but on public trust and legal legitimacy. The deployment of AI-enabled weapons—particularly autonomous systems—raises profound questions about how much lethal authority the public is willing to delegate to machines, and whether current legal frameworks are robust enough to contain that delegation.

International humanitarian law (IHL) and the Geneva Conventions were designed for a world of human combatants, chain-of-command accountability, and physical presence. Autonomous weapons disrupt all three.

Key legal gaps include:

- No agreed definition of "autonomous weapon systems" under international law.
- No clarity on liability when an AI system violates the laws of war.
- No binding treaty regulating the development, deployment, or export of AI-enabled lethal systems.

The UN Convention on Certain Conventional Weapons (CCW) has held multiple rounds of expert discussion on Lethal Autonomous Weapons Systems (LAWS), but progress is stalled. A minority of countries—led by the UK, US, and Russia—continue to resist calls for a binding treaty, citing dual-use concerns and strategic ambiguity.

Meanwhile, public pressure for legal codification grows.

> "The law has not failed AI. The law has simply not caught up." — Christof Heyns, UN Human Rights Council (deceased)

Civil Society and the Push for a Ban

Campaign Leadership

The Campaign to Stop Killer Robots, led by Human Rights Watch and supported by over 70 NGOs, has become a powerful civil society coalition. Their core argument: lethal force must never be fully automated.

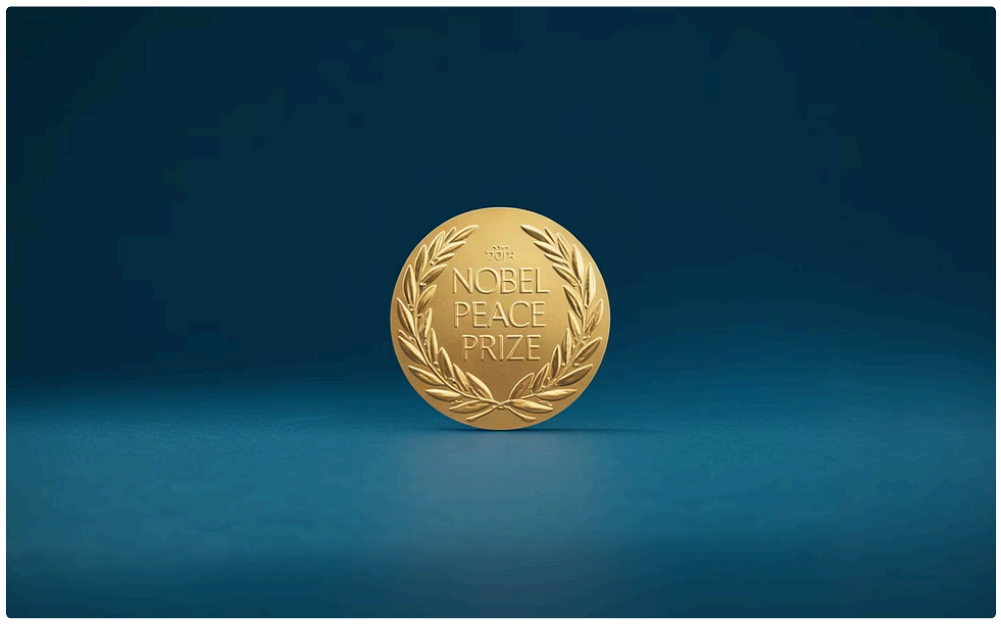

Nobel Laureate Support

Nobel Peace Prize winners like Jody Williams have lent their voices to the campaign, arguing that "The test of our values is not what we deploy in wartime, but what we prohibit in peacetime."

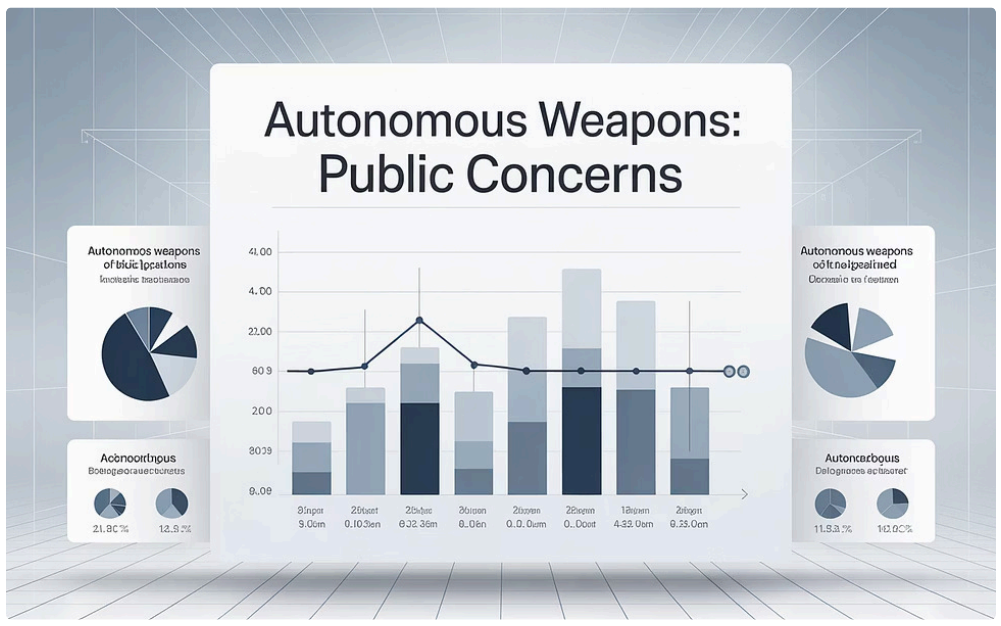

Public Opinion

Recent surveys reveal a clear pattern of public opposition to autonomous weapons. In a 27-country poll, only 27% of respondents supported autonomous weapons without human control. 61% opposed them outright.

Civil society advocates for:

- A pre-emptive, legally binding international ban on fully autonomous weapons
- Clear lines of individual legal responsibility for AI-augmented warfare
- Greater transparency in AI military procurement and deployment

While these groups have no formal role in NATO planning, their influence is real: they shape public opinion, affect national debates in democratic legislatures, and generate reputational risk for defence companies and governments perceived as acting recklessly.

Conclusion: Ethics Is the New Battlespace

The race to integrate artificial intelligence into warfare is not just about speed, scale, or survivability. It is about sovereignty—of command, of judgement, of values. In 2025, the battlefield is no longer bounded by geography alone. It now stretches into code, cognition, and conscience.

As peer adversaries adopt increasingly unrestrained AI doctrines—prioritising tempo over truth, and output over oversight—NATO faces a critical inflection point. To preserve both operational advantage and moral authority, it must act now to embed ethics not merely as a policy overlay, but as a system architecture.

The questions are not abstract:

- Should a machine be allowed to take a human life without a human's consent?
- Who is accountable when an AI system commits an act the world deems a war crime?
- Can public trust be preserved if the electorate does not understand, or agree with, how decisions are made in their name?

If NATO and its partners cannot answer these questions clearly, they risk being outpaced not just technologically, but morally.

Call to Action

Codify a NATO-Wide Ethical AI Standard

Establish shared minimum requirements for human oversight, auditability, and fail-safes—binding across all member state deployments.

Mandate Explainability

Require that all battlefield-deployed AI systems provide actionable, comprehensible explanations to human users in real-time.

Create a Legal Doctrine

Develop a NATO-wide legal framework that clarifies accountability chains, defines thresholds for autonomy, and ensures that liability is never offloaded to the machine.

Engage the Public

Launch transparent, accessible public communications and parliamentary briefings to build informed democratic consent around the use of AI in warfare.

To all those responsible for developing, procuring, deploying, or regulating AI-enabled military systems, this paper recommends these immediate actions, plus:

- Invest in Human–AI Ethical Teaming Technologies: Prioritise funding for systems that embed ethical logic, simulation-based safeguards, and built-in override capabilities—especially at the tactical edge.
- Prepare for Adversarial Legal and Moral Asymmetry: Anticipate how Russia, China, and non-state actors will exploit Western legal gaps and moral hesitations. Incorporate resilience planning into strategy.

> "In the age of intelligent machines, it is not power that will determine who prevails, but principle." — Adapted from General Dwight D. Eisenhower

Ethical AI in warfare is not a restraint. It is a force multiplier—for legitimacy, alliance cohesion, and operational clarity. Nations that lead ethically will not just build better weapons. They will build stronger alliances, command greater public trust, and shape the future rules of war.

The battle for control of AI is already under way. The battle for its conscience has only just begun.