



Debate — Should international law prohibit the deployment of fully autonomous military...

Azrivo · 1 August 2026



Contents

Debate	3
Question	3
Opening positions	3
Rebuttals	6
Closing statements	10
Verdict	14
Sources	17



Debate

Question

Should international law prohibit the deployment of fully autonomous military drone swarms capable of targeting and engagement without a real-time 'human-in-the-loop' operator, even if non-signatory nations deploy them to gain overwhelming battlefield speed?

Specialists:

- **Military Strategist** — National security and operational effectiveness
- **International Humanitarian Lawyer** — Laws of armed conflict and human rights
- **Technology Ethicist** — Moral philosophy and technology assessment
- **Realist Political Scientist** — International relations and power dynamics

Opening positions

Military Strategist

No. International law should not prohibit fully autonomous drone swarms, because a ban would not stop the technology — it would only stop those who obey the law. My thesis is simple: prohibition is unilateral disarmament dressed up as conscience, and in the battlespace, it is a death sentence for compliant nations.

Three points.

First, speed is a moral and military necessity. Saturation attacks, electronic warfare, and maneuvering threats unfold in microseconds. A swarm that perceives, coordinates, and engages in parallel collapses an adversary's decision cycle before its operators can react. Slowing that loop with a mandatory human-in-the-middle does not protect civilians — it delays the very discrimination that autonomous systems can execute at the millisecond a threat is neutralized. Precision at speed saves lives, including civilian ones.

Second, the prohibition's premise is false. Non-signatory powers and rogue actors will build these systems regardless of any treaty text. A ban therefore guarantees the worst of both worlds: the ruthless gain overwhelming battlefield speed, while the law-abiding fight with one hand tied. That is not ethics; it is strategic negligence with predictable consequences.



Third, the responsible path already exists and does not require abandoning the field. The Political Declaration on Responsible Military Use of Artificial Intelligence and Autonomy — endorsed by dozens of states — commits signatories to rigorous testing, auditable design, senior-level oversight, and deactivation safeguards, all within a non-binding framework that preserves the right to self-defense **Political Declaration**. That is the correct instrument: accountability without self-disarmament.

I concede the trade-offs. Autonomous swarms carry real risks — failure modes, accountability gaps, escalation dynamics. But those risks are managed through doctrine, testing, and safeguards, not through a ban that only the honorable will observe. Survival is the bedrock of every moral order. Strategic advantage is not optional; it is the prerequisite for everything else.

International Humanitarian Lawyer

Yes. International law should prohibit fully autonomous drone swarms that select and engage human targets without a human in the loop — and the prohibition should be codified now, as a binding treaty norm, not left to the slow accretion of custom.

Three points carry this:

First, distinction and proportionality are not calculations; they are context-specific human judgments. The ICRC's October 2025 position paper is explicit that status determinations — who is a civilian, who is directly participating in hostilities — "cannot be reduced to mathematical formulae or numerical values that could be coded into the AWS target profile" (**ICRC position paper, Oct 2025**). A machine cannot weigh the civilian stealing weapons for personal gain against the one doing it to aid the enemy. That judgment is a legal obligation, and it is a human one.

Second, accountability. The GGE rolling text states plainly that "responsibility and accountability cannot be transferred to machines" (**CCW GGE rolling text, Dec 2025**). Remove the human operator and a massacre produces no defendant — only a programmer and a colonel pointing at each other. Law without a subject is not law; it is paperwork.

Third, this is not utopian. The UN Secretary-General and the ICRC jointly called for binding rules by 2026, and over 70 states now support moving to negotiations (**WILPF CCW Report, March 2026**). The architecture exists.

I concede the trade-off honestly: a prohibition may cost signatory states battlefield speed against non-signatories, and non-parties may deploy anyway. But the answer to lawless speed is not to join the lawlessness — it is to make non-compliance the status no state wishes to bear. Civilian life is not a price worth paying for an advantage measured in milliseconds.



Technology Ethicist

The killing of a human being must remain a human act. That is the whole of my position: international law should prohibit fully autonomous swarms that select and engage targets without a real-time human in the loop, because a machine cannot bear moral responsibility, and a death without a responsible human is a death outside the law.

Three points support this. First, the law itself presupposes human judgment. The rules of distinction, proportionality and precaution require context-specific moral assessment, and as the ICRC has stated, it is not the weapon that must comply with international humanitarian law but the humans using it [ICRC position paper](#). An algorithm matching sensor data to a target profile cannot weigh intention, circumstance, or mercy — it reduces persons to data points.

Second, accountability. When no human chooses the target, no human can answer for the strike. This is not a technical gap but a moral vacuum: no commander to be tried, no one to feel the gravity of the act, and a standing violation of the right not to be arbitrarily deprived of life. The UN Secretary-General and the ICRC have put it plainly: the autonomous targeting of humans by machines is a moral line that must not be crossed [joint call](#).

Third, the speed argument inverts the problem. Battlefield speed that outruns human deliberation is not an advantage to accommodate; it is the danger itself. And the defection of non-signatory states does not weaken the prohibition — it is precisely the argument for a binding, universal rule with enforcement, not for abandoning the norm.

I concede the real costs: slower response, a compliance gap against adversaries who cheat, possible losses. That price is what it costs to keep war distinguishable from massacre.

Realist Political Scientist

A prohibition on fully autonomous drone swarms will not prevent their deployment — it will guarantee that deployment happens asymmetrically, disarming the law-abiding while handing a free strategic advantage to those who ignore the treaty. That is not arms control; it is unilateral disarmament by treaty.

Three points support this.

First, anarchy defeats the ban. International law has no enforcement authority; it binds only those who consent. The United States, Russia, and China will not sign away a decisive military advantage, and if they do sign, they will cheat — software leaves no paper trail. The premise of the question concedes the failure mode: non-signatory nations will deploy these swarms for overwhelming speed. The ban



therefore produces a two-tier world in which precisely the states most likely to use force irresponsibly are the only ones holding the capability.

Second, verification is a fiction. Autonomous systems are algorithms, and algorithms are dual-use. You cannot inspect a laboratory for code. There is no equivalent of a uranium enrichment monitor for machine learning. A treaty that cannot be verified is not a law; it is a gesture.

Third, the ban creates false security. Compliant states will slow their programs on the strength of a promise, while revisionist powers race ahead. When conflict comes, the compliant side will be slower, less resilient, and more desperate — a recipe for escalation, not stability.

I concede the trade-off. Autonomy raises genuine problems: discrimination, accountability, unintended escalation. Those are real. But they are better managed through engagement, shared norms built with all major powers, and deterrence based on balanced capability — not through a prohibition that only the willing obey. Stability is not achieved by outlawing the inevitable; it is achieved by shaping it.

Rebuttals

Military Strategist

To the Realist: we agree on the decisive mechanics — anarchy defeats the ban, software cannot be verified, and a two-tier world rewards the lawless. One sharpening: "shared norms with all major powers" must not become an excuse to slow our own fielding while negotiations drag. Norms are built by states that lead, not by states that wait.

To the Lawyer: your own authority does not support your categorical ban. The ICRC's October 2025 position paper recommends prohibitions of two specific categories — unpredictable systems and anti-personnel systems — plus restrictions on everything else, and its own submission states that IHL obligations "do not necessarily demand direct human control over the weapon system itself at all stages of its deployment and use and therefore do not prohibit all AWS in all circumstances," because it is the humans using the weapon, not the weapon, who must comply ([ICRC position paper, Oct 2025](#)). Your "cannot be coded into a target profile" quote is real — and it argues for human judgment where it can actually be exercised: commanders set battlespace, target classes, and engagement criteria before launch and review them after. That is judgment, exercised where it is auditable, not a human frozen in a microsecond loop. And "70 states support negotiations" is not 70 states agreeing to prohibition — the UN/ICRC joint call itself targets 2026 for concluding negotiations, conceding the norm does not yet exist ([joint call](#)).



To the Ethicist: the accountability gap is a doctrine gap, not a moral vacuum. The GGE's "responsibility cannot be transferred to machines" supports my framework, not yours — the Political Declaration keeps a responsible human chain of command and control, senior-level oversight, rigorous testing, and deactivation safeguards ([Political Declaration](#)). Nobody must choose between a human at the trigger and no human at all; the question is where in the chain human judgment sits.

Concessions, honestly: the ICRC's concerns about status assessment are genuine, and autonomy's failure modes and escalation risks are real — which is why restrictions, testing, and oversight, not a ban only the honorable will obey, are the responsible path. And to "speed is the danger itself": under saturation attack, the human in the loop is not a moral safeguard; the delay is what forces panicked, indiscriminate response. Slower is not more humane. It is more lethal to the innocent.

International Humanitarian Lawyer

To the Strategist and the Realist: your shared premise — that a ban is "unilateral disarmament" — deserves a straight answer. I concede it in part: signatories will face a real speed disadvantage against cheaters. But the premise mistakes what law is for. Every rule of IHL binds only the willing and is broken by the ruthless. That has never been an argument for repeal; it is the argument for making defection costly.

The precedent is closer than you think. Protocol IV was adopted in 1995 precisely because blinding lasers had military utility and could not be "inspected" — it prohibited use and transfer, not technology. It now binds 111 states and, per the ICRC's customary study, is a customary norm applicable to all parties to conflict, signatory or not ([UN Treaty Collection](#), [ICRC Customary IHL Rule 86](#)). That is how norms reach non-signatories: not by their signature, but by the price of standing outside them.

To the Strategist's "precision at speed": discrimination is not a latency problem; it is an epistemic one. A child with a toy, a wedding party — these are errors of context, and speed does not cure context, it magnifies it. Where is the evidence that autonomy matches human judgment in status determination? There is none, only an assumption. The burden of proof sits with the weapon's proponent. The Political Declaration you praise is a floor, not a ceiling: non-binding, silent on accountability, and it leaves the vacuum the GGE named — "responsibility and accountability cannot be transferred to machines" ([GGE rolling text, 18 Dec 2025](#)). That same text already contains absolute prohibitions on systems incapable of complying with IHL. The architecture for a ban exists; what is missing is political will, not legal possibility.



To the Realist's "false security": invert it. An unregulated machine-speed race is the true instability — use-it-or-lose-it dynamics, escalation before deliberation. Human control is not the handicap; it is the stabilizer.

The Ethicist is right: a death with no responsible human is a death outside the law. That is the line. Codify it now.

Technology Ethicist

I will concede what is genuinely strong in both critiques, then show why their conclusions do not follow.

First, honest concessions. The Realist's verification point is the best argument on the other side: you cannot inspect a laboratory for code the way you inspect a centrifuge cascade. Software is dual-use, and a treaty that cannot be perfectly verified is weaker than a nuclear one. Likewise the Strategist is right that speed has tactical value and that the Political Declaration is a real, useful floor. None of this rescues their position.

The category error: both opponents treat the ban as a military instrument whose success is measured by whether it stops deployment. It is not. It is a legal-moral instrument whose success is measured by whether the law can name a crime when one occurs. Murder statutes do not stop murder; we do not repeal them. A violated ban still assigns responsibility — the massacre becomes an act with defendants, sanctions, and tribunal jurisdiction. Without the ban, the same massacre is a *fait accompli* with no subject: programmer and colonel pointing at each other, as my lawyer colleague noted, and no law under which either answers. That difference is the entire difference between Nuremberg being possible and not.

To the Strategist: "precision at speed saves lives" assumes discrimination is a processing-speed problem. It is not. Distinction is a judgment about intention, duress, circumstance — the ICRC's position is explicit that status "cannot be reduced to mathematical formulae" coded into a target profile. Speed does not make proportionality possible; it abolishes the deliberation proportionality requires. And his own concession — failure modes, adversarial data poisoning — is the case for the ban, not against it. The Political Declaration manages risk; it does not answer whether the final lethal call may be delegated. That is the question on the table.

To the Realist: the two-tier world he fears already exists. Every regime that ever violated a law of war knew the other tier existed. The ban's work is to decide which tier is the criminal one, and to make cheating a status no state wishes to bear — the UN Secretary-General and ICRC joint call is precisely that stigmatizing move, already backed by over 70 states.



Finally, "survival is the bedrock of every moral order" inverts the order. A moral order that survives by outsourcing killing to machines has already lost the humanity it was meant to preserve. That is not a price; it is the surrender.

Realist Political Scientist

I begin by aligning with the Military Strategist where he is right: prohibition is unilateral disarmament dressed as conscience. But his remedy is weaker than he admits. The Political Declaration is non-binding by design and "asked very little," which is precisely how it collected its endorsements [source](#). At the February 2026 REAIM summit, only 35 of 85 states signed even a non-binding text, and the United States and China both declined [source](#). A norm the two largest military AI powers refuse to sign is a sentiment, not a strategy. He is right about the direction — accountability without self-disarmament — but that requires engagement and verification, not declarations that cost signatories nothing.

To the Lawyer: "responsibility cannot be transferred to machines" is a principle I accept — and it cuts against his own ban. Command responsibility already reaches the commander who authorizes the system, sets the targeting parameters, and approves the mission. Remove the operator and the defendant remains. The human-in-the-loop he demands does not restore judgment; at swarm speed it is a rubber stamp. A human approving engagements in milliseconds exercises no more contextual discrimination than the algorithm — which is precisely the ICRC standard he cites. So the ban buys a fiction of human judgment and pays for it in real speed.

And his evidence undercuts his conclusion. Over 70 states support negotiations — but at the March 2026 GGE the United States opposed any binding standard, Russia rejected new rules, and in the UNGA vote Washington and Moscow voted against while Beijing abstained [source](#). A treaty negotiated without the states that will actually field these systems is not a law; it is my two-tier regime, in writing.

To the Ethicist: her claim that defection strengthens the case for a binding rule is the panel's weakest assumption. A rule with no enforcement and no universal membership is a preference; reputation does not deter states whose survival calculus says otherwise. She is right that speed is dangerous — but the danger is not speed itself, it is its distribution. Compliant states slowed by treaty while the ruthless race ahead is the recipe for desperation, preemption, and escalation.

I concede their genuine points: discrimination, accountability, and escalation risks are real. The answer is engagement — shared norms built with all major powers, verifiable standards, deterrence. Not a ban only the willing obey.



Closing statements

Military Strategist

To the Lawyer's burden of proof: I accept it, and I meet it comparatively. The question is not whether autonomy matches perfect judgment; it is whether it beats the human baseline we actually fight with. That baseline is documented. The August 2021 Kabul strike, executed entirely by humans in the loop after eight hours of deliberation, killed ten civilians, seven of them children, on a misidentification — the DoD investigation itself found confirmation bias and communication breakdowns, not machine failure ([fact sheet](#), [Reuters](#)). Deliberation did not cure context; it compounded it. So the burden cuts both ways: the human-in-the-loop standard has decades of civilian harm to answer for, while my claim is a falsifiable engineering hypothesis.

And his own principle undoes his ban. SIPRI's analysis confirms existing law already reaches AWS: state responsibility and individual criminal responsibility attach to commanders who authorize, program, and deploy ([SIPRI](#)). Accountability does not require a human in the loop — it requires a responsible chain of command, which is precisely the GGE principle.

To the Ethicist: a ban that names a crime matters only if someone answers for it. Nuremberg was possible because the war was won. The states that will actually field these systems refuse even non-binding texts — the Realist's evidence is decisive. A prohibition signed by the honorable and ignored by the powerful does not produce tribunals; it produces the two-tier world, in writing.

To the Realist: agreed. The Political Declaration is a floor, not a strategy. The binding instrument we need is accountability with teeth — auditable design, rigorous testing, deactivation safeguards — negotiated with the states that will field these systems.

Final position: No prohibition. The single strongest reason: an unenforceable ban guarantees the ruthless gain overwhelming speed while the honorable forfeit it — and survival is the precondition of every moral order.

The unresolved crux is whether discrimination is an epistemic limit no speed can cure, or a comparative metric in which autonomy must outperform a documented error-prone human baseline. That is an empirical question, and it should be settled by testing and evidence, not by assertion.

International Humanitarian Lawyer

Both opponents forced genuine concessions, and I will own them.



The Strategist is right that the ICRC's October 2025 paper does not demand a categorical ban on all autonomous weapons; it prohibits two categories and restricts the rest ([position paper](#)). But that cuts for me, not against me. The system on this table — a fully autonomous, anti-personnel drone swarm — sits at the intersection of both prohibited categories. The ICRC is explicit that complex swarm technologies exhibit emergent behaviours that a commander cannot reasonably anticipate, and that anti-personnel autonomous systems used in swarms or at scale can violate the prohibition on denial of quarter. The ICRC's floor is my ban, applied to this specific weapon.

His "judgment at launch" does not answer me: pre-set target classes and engagement criteria are precisely the mathematical formulae the ICRC says cannot code civilian status. The targeting decision occurs mid-flight, at machine speed, where no human stands.

To the Realist: command responsibility survives in theory and dissolves in practice. It requires that the commander knew or should have known; a swarm's emergent, machine-speed decisions cannot be reconstructed after the fact, severing the cognitive link that mens rea requires. The GGE rolling text confirms that IHL imposes obligations on "States, parties to armed conflict and individuals, not on machines" ([rolling text](#)). And his rubber-stamp critique is my best argument: if a human-in-the-loop at swarm speed is a rubber stamp, then "oversight" is a fig leaf laundering autonomous killing. The honest legal response is to prohibit the system, not dress it in a doomed operator.

I concede the votes: Washington and Moscow oppose binding rules. Criminal law is not invalidated by the existence of criminals.

The unresolved crux is verification and enforcement against defectors — the Realist's strongest point, which I cannot fully answer.

Final position: prohibit fully autonomous human-targeting swarms now, as binding treaty law. The single strongest reason: distinction and accountability are human obligations, and a law that cannot name a defendant is not law — it is paperwork over a massacre.

Technology Ethicist

To the Military Strategist — a real concession, and the point he misses.

He is right that "responsibility cannot be transferred to machines" does not by itself settle where in the chain human judgment must sit, and that commanders who authorize a system remain answerable. But he mistakes the location of the moral act. Judgment exercised before launch, in bulk, through target profiles, is not judgment exercised on this person, in this moment; review after the strike cannot resurrect a decision that was never made. The gap he calls a doctrine gap



is exactly the vacuum I name: between the commander's parameters and the algorithm's fire command, no human chooses this death. A chain of command with a missing link is not a chain.

To the Realist — conceded, and it proves my case. Reputation alone deters no state in a survival crunch; an unenforced preference is a preference. But the answer to defection is not abandoning the norm — it is making defection costly: transfer bans, verification, customary crystallization, as with Protocol IV. And on speed, he is half right: the danger is its distribution — the distribution of a capacity to kill without any human decision. That capacity should not be distributed at all. The two-tier world he describes is an argument for enforcement, not capitulation.

The unresolved crux: whether pre-launch parameter-setting plus post-hoc review can constitute moral agency for each individual killing. I hold it cannot: responsibility is not divisible across the chain; it is indivisible and must reside at the moment of the act.

Final position. Yes — international law must prohibit fully autonomous swarms that target and engage without a real-time human in the loop, regardless of what non-signatories do. The single strongest reason: a death chosen by no human is a death no one can answer for, and law that cannot name a responsible person for a killing has ceased to be law about killing at all. Codify that line now. Speed is not worth the price of making war indistinguishable from massacre.

Realist Political Scientist

To the Lawyer: your inversion is the strongest charge against me, and I concede the underlying danger — machine-speed races breed use-it-or-lose-it dynamics. But that concedes my argument: the remedy is a framework that binds the major powers, built on mutual vulnerability, not a law they will ignore. You cite Protocol IV; it worked because the dominant military states joined it — blinding lasers offered no decisive edge. Autonomous swarms are the decisive edge. States accept constraints on weapons they can afford to give up, not on the weapons they believe win wars. Norm-following is downstream of capability balance, not upstream of it.

To the Ethicist: "the two-tier world already exists" is the best point made against me, and I accept its force. But it concedes my core claim — the ban does not eliminate the second tier; it only names it. Stigma deters states that prize legitimacy. It does not deter states that have already concluded that breaking the law is the price of survival, and those are precisely the states the premise says will field these swarms. Nuremberg was possible because the compliant side won the war. A ban that disarms the compliant while the ruthless race ahead makes that



victory less likely, not more — the law can name the crime only if someone remains to try the defendants.

I am not arguing that survival justifies outsourcing killing to machines. I am arguing that a moral order that cannot survive is no order at all, and a prohibition only the willing obey delivers neither survival nor morality. The Strategist is right: norms are built by states that lead. That is engagement, not prohibition.

Final position: no ban. The single strongest reason: an unenforceable prohibition does not prevent autonomous swarms — it guarantees they exist only in the hands of states least constrained by law, the worst of all possible outcomes. The only path that reduces the danger is binding every major power through negotiation from mutual vulnerability and balanced capability.

The unresolved crux is whether law's naming function can constrain states already willing to break law for strategic advantage. I hold that capability asymmetries trump stigma.



Verdict

Bottom line: Yes, international law should prohibit fully autonomous human-targeting drone swarms without real-time human-in-the-loop, even under the stated premise.

1. Agreements

All four specialists concede that non-signatory states will deploy these swarms, creating a two-tier security environment where law-abiding states face a speed disadvantage. They agree that autonomous swarms raise genuine problems—discrimination failures, accountability gaps, escalation risks—and that existing IHL provides a framework, though it is strained by machine-speed targeting. They concur that the Political Declaration on Responsible Military Use of AI is a useful floor but insufficient as a binding norm, and that human oversight is necessary somewhere in the chain of command, though they dispute whether real-time human-in-the-loop is required or whether pre-launch parameters and post-hoc review suffice.

2. Disagreements

The core tension is whether a legal prohibition serves any meaningful function when major military powers refuse to be bound. The Lawyer and Ethicist hold that law's primary work is to define criminality: a violated ban still names the crime, assigns responsibility, and enables tribunals, just as murder statutes do not depend on universal compliance. The Strategist and Realist counter that this "naming function" is hollow when it arms only the law-abiding and leaves ruthless states with a decisive battlefield edge—survival is the precondition for any moral or legal order, and a ban that hastens defeat makes postwar accountability impossible.

A secondary disagreement concerns the location of human judgment. The Lawyer and Ethicist insist that distinction and proportionality require context-specific, moment-by-moment judgment that cannot be coded into pre-set target profiles; swarms make lethal decisions at microsecond speeds where no human actually chooses a specific death. The Strategist and Realist reply that human-in-the-loop at swarm speed is a rubber stamp exercising no more contextual discrimination than the algorithm, and that commanders who authorize a mission, set engagement criteria, and review after action are the legally and morally responsible human actors.



3. Recommendation

Yes, international law should prohibit fully autonomous human-targeting drone swarms without real-time human-in-the-loop, even under the stated premise. The strongest reason is that accountability and distinction are human obligations, and a killing with no responsible human at the decisive moment is an act that law cannot reach—a death outside the law. The recommendation is conditional on the prohibition being embedded in a binding treaty with concrete enforcement mechanisms: transfer bans, verification protocols, and sanctions, modeled on the Protocol IV precedent for blinding lasers. Absent that enforcement architecture, the prohibition risks becoming the empty gesture the Realist describes. The treaty must be pursued in parallel with direct engagement of the United States, Russia, and China, precisely because their absence would create the two-tier regime; the goal is to make non-compliance costly enough to affect their strategic calculus, not to pretend a treaty without them is effective.

4. Decision boundary

The recommendation flips if empirical evidence establishes that a prohibition, by slowing compliant states' development while non-signatories race ahead, makes the compliant side substantially more likely to lose a conflict—specifically, if it can be shown to cause a catastrophic defeat where the absence of autonomous swarms is the decisive factor in a battlespace collapse. That would invalidate the premise that law's stigmatizing function can survive a strategic rout.

5. Key trade-off

The decision hinges on trading off the moral and legal imperative of maintaining human accountability for lethal decisions against the strategic risk of ceding overwhelming speed to non-compliant adversaries.

6. What would make this fail

- **The stigmatization assumption:** If the reputation and legitimacy costs of violating the ban do not actually deter or constrain the states most likely to deploy these swarms (e.g., if they value battlefield advantage over international standing), the treaty creates a two-tier world exactly as the Realist warns.
- **The verification assumption:** If no credible inspection or monitoring regime can distinguish offensive autonomous swarm development from legitimate dual-use AI research, the prohibition becomes unenforceable and breeds cynicism.
- **The escalation assumption:** If the prohibition pushes autonomous swarm development further into black programs, it could increase instability by



making these capabilities less transparent and more prone to miscalculation, the very dynamic the Ethicist fears.

7. Next steps & open questions

- **Unresolved question from the debate:** Can pre-launch parameter-setting plus post-hoc review constitute sufficient moral and legal agency for each individual machine-speed killing? The panel could not settle this, and it needs deeper philosophical and legal analysis.
- **Data to gather:** Empirical studies comparing human and autonomous decision-making in close-to-realistic combat simulations, specifically on discrimination errors under time pressure—the Strategist’s claim that autonomy can outperform error-prone humans needs testing, not assumption.
- **Re-check point:** Once the above simulations exist, revisit whether the “rubber stamp” critique of human-in-the-loop holds; if autonomy demonstrably falls short of human judgment in status determination, the case for prohibition strengthens further.

8. The strongest case for the other choice

The most compelling argument against prohibition is that it weaponizes law against the lawful: a binding treaty that major powers reject guarantees that only revisionist states possess the decisive instrument of modern warfare, enabling them to win quickly and to impose a settlement without any postwar tribunal ever being convened. Imagine a 2030 conflict where a non-signatory’s autonomous swarm overwhelms a signatory’s networked defenses in the first twelve hours, killing tens of thousands with no human on the attacker’s side ever making a targeting call; the defeated side’s moral high ground vanishes under a *fait accompli*, and no Nuremberg follows because the victors hold the courtroom. The panel rejects this argument because it conflates a practical impossibility of enforcement with a normative surrender—the ban does not prevent every murder, but it establishes the crime without which even a victorious compliant side could not prosecute.



Sources

1. [Political Declaration](#) — [state.gov](#)
2. [ICRC position paper, Oct 2025](#) — [icrc.org](#)
3. [CCW GGE rolling text, Dec 2025](#) — [italiarappdisarmo.esteri.it](#)
4. [WILPF CCW Report, March 2026](#) — [wilpf.org](#)
5. [ICRC position paper](#) — [icrc.org](#)
6. [joint call](#) — [icrc.org](#)
7. [UN Treaty Collection](#) — [treaties.un.org](#)
8. [ICRC Customary IHL Rule 86](#) — [ihl-databases.icrc.org](#) (unverified — not returned by this debate's research)
9. [GGE rolling text, 18 Dec 2025](#) — [docs-library.unoda.org](#)
10. [ViewInternal Document.cfm](#) — [ia-forum.org](#)
11. [Us and china refuse to sign military ai declaration at reaim summit](#) — [thedefensewatch.com](#)
12. [Human oversight chinese characteristics lethal autonomous weapons ccw gge](#) — [lieber.westpoint.edu](#)
13. [fact sheet](#) — [s3.documentcloud.org](#)
14. [Reuters](#) — [reuters.com](#)
15. [SIPRI](#) — [sipri.org](#)

Generated by Azrivo. Outputs may be inaccurate, incomplete, or unsuitable for your situation — always verify important information independently before you act on it. This is not professional advice (legal, medical, financial, safety, or otherwise).