State Responsibility in Relation to Military Applications of Artificial Intelligence: An Article Summary and Analysis

Adeyemi Folarin

SRT751NBB

Professor Masood Erfanian

July 2, 2026

Table of Contents

Table of Contents 2
Introduction 3
Main Points 4
- **Military AI Is Bigger Than Autonomous Weapons** 4
- **Attribution: Finding the Human in the Loop, or Behind It** 4
- **Compliance-by-Design Before Anything Is Fired** 5
- **Responsibility for What Others Do** 6

The Article's Conclusion 7
Feedback and Thoughts 8
Information Sources 9

Introduction

Bérénice Boutin's article "State Responsibility in Relation to Military Applications of Artificial Intelligence," published in the Leiden Journal of International Law in 2023, deals with a problem that has come up again and again in this course: technology moving faster than the laws that were written before it existed. Militaries now use AI for surveillance, logistics, target recommendation, and battlefield decision support, but most of the legal writing has focused on one narrow slice of that picture, the "killer robot." Boutin, a senior researcher at the Asser Institute in The Hague, takes a wider view and asks what international law already has to say when these systems contribute to a violation.

Her research question is: under what conditions can a state be held responsible under international law for violations involving military AI? Her approach is to work with the law that already exists. Rather than proposing a new treaty, she tests whether the law of state responsibility, written down in the International Law Commission's Articles on Responsibility of States for Internationally Wrongful Acts (ARSIWA), can handle the unusual properties of AI: autonomy, opacity, unpredictability, and machine speed. She organizes the analysis around three stages of the technology's life. First, deployment, where the problem is attributing wrongful conduct to the state. Second, development and procurement, where she argues for a duty she calls compliance-by-design. Third, derived responsibility, where a state can be held responsible because of its links to what other states and private companies do.

Her main argument through all three stages is that there is no responsibility gap at the state level. The law can still reach these systems; what changes is where you have to look for the relevant human decision.

Main Points

Military AI Is Bigger Than Autonomous Weapons

Boutin starts by pushing back on the assumption that military AI means weapons that select and engage targets on their own. She distinguishes rule-based systems, where a programmer encodes every instruction, from machine learning systems trained on data, and shows why the machine learning kind causes legal problems: their outputs are hard to explain (the black box problem), their behaviour is hard to predict, and they operate at a speed no human can keep up with in real time. That speed leads to automation bias, where an operator who is formally supposed to approve a system's recommendations has no real ability to evaluate them and rubber-stamps them instead.

Her examples make this concrete. Project Maven, a US Department of Defense program, uses machine learning on drone footage to save analysts from repetitive video review and feed decision-making. Project Convergence goes further, networking drones, digital maps, threat-prioritization algorithms, and fire control systems so ground operators can act on sensor data within seconds. Neither is a weapon in the traditional sense, yet both sit inside a kill chain. Data from a system built for surveillance can end up in a targeting process, which is why Boutin argues that any legal line between lethal and non-lethal AI falls apart in practice.

Attribution: Finding the Human in the Loop, or Behind It

The law of state responsibility connects the acts of people to the state they serve. Boutin points out the tension AI creates: attribution assumes there is human conduct causally linked to the violation, and an autonomous system can stretch that link until it breaks. Some scholars think

attribution is easy here since anything a state organ does in official capacity counts as an act of the state.

Boutin disagrees. The first question is whether the operator's involvement even counts as human conduct when the machine, not the person, drove the outcome.

Her answer is a framework with four scenarios. Where an operator has real control, such as an object-recognition aid used alongside the operator's own eyes, attribution works through the operator's conduct as usual. Where a system runs on its own once activated, like an air defence battery in automatic mode, the meaningful human decision moves up the chain to the commanders and officials who chose to deploy it and set its parameters, so attribution runs through their conduct under Article 4 ARSIWA. The grey zone in between, decision-support tools with a human veto, gets her most careful treatment: because machine speed can make that veto meaningless, attribution may again fall on the people who made the deployment decisions rather than the operator. For hypothetical future systems so autonomous that no human conduct connects to the harm, she admits you could treat the AI itself as a state agent under Article 8 ARSIWA, but she rejects that option because it quietly accepts the idea that software has agency or legal personality. The example she uses to illustrate the middle scenario is the targeting-support system that gathers sensor data and proposes targets while a human keeps the launch decision in name only.

Compliance-by-Design Before Anything Is Fired

The second pillar builds on a feature of state responsibility that most public debate ignores: a violation does not require any damage to have happened. A state violates international law the moment it breaks a binding obligation, which means responsibility can kick in while a system is still in the lab or the procurement pipeline. Boutin walks through the obligations that apply at

this stage. Common Article 1 of the Geneva Conventions requires states to ensure respect for international humanitarian law, which she interprets as a duty to design and train algorithms to meet IHL standards and to drop systems that cannot.

Article 36 of Additional Protocol I requires legal review of new weapons, means, and methods of warfare, and she argues AI decision-support tools fall within "means or methods" even when they are not weapons. Human rights obligations cover dual-use systems built outside any military program.

Her strongest claim in this section is about procurement. A state that buys AI from a contractor cannot hide behind the vendor's trade secrets or the model's complexity. Wilful blindness is not a defence; the state has to demand information, test the system, and verify compliance before using it. She also raises an open question the framework cannot settle on its own: whether negligence is enough for technologies this unpredictable, or whether states should negotiate a strict liability regime, and she points to the Convention on Certain Conventional Weapons meetings as the place for that debate.

Responsibility for What Others Do

The third pillar covers conduct one step removed. Under Article 16 ARSIWA, a state that knowingly aids another state's wrongful act shares responsibility. Boutin's coalition example makes the stakes concrete: a state that feeds AI-generated target lists to a partner becomes responsible once it learns the partner keeps striking those targets without proper human verification. Arms exports raise the same kind of duty, and she connects AI to the Arms Trade Treaty, whose Articles 6 and 7 prohibit transfers a state knows will be used for war crimes and require risk assessment for the rest.

Private industry gets similar treatment. States must regulate companies within their jurisdiction whose AI could lead to human rights violations, a duty she bases on the ICCPR and on Council of Europe work concluding that voluntary AI ethics pledges from the tech sector are no substitute for enforceable law. Since these firms operate across borders, she argues, no state has an excuse to pass the buck.

The Article's Conclusion

Boutin closes by fitting the three pillars together into one argument that covers the technology's whole lifecycle. Attribution covers the battlefield, compliance-by-design covers the lab and the procurement office, and derived responsibility covers allies and vendors. Together, they support her thesis that state responsibility leaves no gap for military AI to fall through: at every stage, there is some human decision attributable to the state that can serve as the hook for legal responsibility, even when the operator's finger on the trigger no longer matters.

She explains why this matters in governance terms. Because states build, buy, regulate, and deploy these systems, holding states responsible reaches the bigger structural forces behind military AI in a way that prosecuting individuals cannot, and the two systems work side by side rather than competing. This approach also works as prevention: the duty to build legal norms in at the design stage turns state responsibility from an after-the-fact liability rule into a tool for shaping how engineers build the technology.

She ends with an honest concession. The framework exists on paper, but putting it into practice requires interdisciplinary work the law cannot do alone: engineering methods, testing and certification protocols, and policy guidance that turn legal principles into technical standards. She lays out that research agenda instead of pretending the law will enforce itself.

Feedback and Thoughts

What convinced me most is how the attribution framework moves responsibility up the chain. Instead of pinning fault on the soldier staring at a screen whose output he has no real ability to question, Boutin follows the causal chain up to the procurement officer and the commander who set the parameters. That matches how accountability should work in any complex system. In security operations we see the same dynamic with alert triage: the real decision points sit with the people who tuned the detection logic and set the auto-response thresholds, not with the analyst drowning in SIEM alerts who clicks approve. Boutin turns that intuition into an actual legal argument.

Compliance-by-design also lines up with things this course and my program cover under a different name. It is security-by-design and shift-left thinking carried over into international law: build the constraint into the system before deployment instead of fighting about the failure in court afterward, and treat "we could not have known" as negligence. Her refusal to accept vendor secrecy as an excuse at procurement is the legal cousin of demanding a software bill of materials before you deploy a product on your network.

Two weaknesses stand out to me. First, she admits principles like proportionality may be impossible to turn into code, and her fallback, stop development or add human involvement, runs into her own evidence that human oversight of fast AI is often just for show. The article identifies the automation bias problem and then leans on human involvement as the remedy anyway. Second, the framework has the same enforcement weakness as international law in general: the states most invested in military AI are the least likely to accept any tribunal applying ARSIWA against them. Boutin talks about putting the framework into practice, but there is still no enforcement mechanism with real teeth.

Even so, the article does what this course asks us to look for in the relationship between law and technology: it shows existing law stretching to cover a new technology, and it shows the technology forcing the law to look for its oldest concept, the responsible human, in new places in the chain of command. The common claim that AI is unregulated turns out to be lazy. The regulation exists; states would rather not apply it to themselves.

Information Sources

Boutin, Bérénice. "State Responsibility in Relation to Military Applications of Artificial Intelligence." Leiden Journal of International Law, vol. 36, no. 1, 2023, pp. 133-150, https://doi.org/10.1017/S0922156522000607.

"State Responsibility in Relation to Military Applications of Artificial Intelligence." Cambridge Core, Cambridge University Press, 28 Nov. 2022, https://www.cambridge.org/core/journals/leiden-journal-of-international-law/article/state-responsibility-in-relation-to-military-applications-of-artificial-intelligence/1B0454611EA1F11A8B03A5D2D052C2BE.