
Artificial Intelligence as a Weapon of War

Open Access Teaching Case Developed for the Tech for Humanity Pathways Minor

Funded by the Andrew Mellon Foundation

Developed by Elena Roe

Background

In recent years, talk of the use of Artificial Intelligence (AI) across numerous fields has exploded, and the realm of the military and warfighting is no exception. An increased weaponization of AI by numerous military superpowers has been conceptualized as a new arms race, with the Cold War escalation of nuclear capability substituted by automated weapons systems (AWS) (Humble, 2024). The rise of AI in the international community can be traced back to 2012, when the U.S. Department of Defense (DoD) released directives regarding the use of AWS; the same year, Human Rights Watch and the Harvard Law School released a joint report calling for the outright prohibition of the use of AWS under international law. In the years since, minimal progress has been made in the realm of international law and AI in conflict. There is currently no legal framework of international regulatory bodies that outlines ethical use of AI, especially in the context of war, leaving nations and their militaries essentially free to develop—and engage in the AWS arms race—as they see fit.

To backtrack, however, what precisely is an AWS, and where/how are they currently being used? In 2023, the U.S. DoD defined autonomous weapons systems as “being capable of, once activated, to select and engage targets without further intervention from a human operator” (Department of Defense 2023). Examples of AWS include Israel’s Iron Dome air defense system, the German MANTIS, and Sweden’s LEDS-150; further, the South Korean Super aEgis II qualifies as an AWS, but it is used for surveillance during peacetime (Humble 2024). These kinds of weapons have also been used with increased frequency in the Russo-Ukrainian war, where autonomous drones—notably needing no human interference in order to identify and

strike targets—have been employed in a major conflict for the first time (Marr, 2024). These “killer robots” have been used by the Ukrainian army alongside autonomous robots such as the dog-like “BAD One”, which can detect enemy placements, mines, and even carry ammunition to soldiers on the battlefield; it is also equipped with autonomous machine guns. Ukraine has also employed AI to monitor supply routes, resettle refugees, and surveil troop movements, while also utilizing facial recognition technology in an attempt to track war criminals’ activity and hold them to account.

The increased use of AI in warfare has, however, raised serious ethical and logistical questions as to the participation of AI developers in conflict scenarios, given the intention of “harmless” developments of this kind of technology. The first concern that typically arises concerns the limitations of AWS, as the systems are vulnerable to poisoning and evasion (Starck et. al 2022), and have been shown to be inaccurate in many contexts. For example, AI has been shown to be unable to recognize individuals with darker skin tones, which exemplifies the potential for fatal inaccuracies when facial recognition software is utilized in war—horrifically, to the extent that individuals are targeted to be killed (Stewart and Hinds, 2023). Beyond AI’s limitations, the question remains: should a machine ever be given free rein to kill? With all of the moral ambiguities and ethical decisions present in warfighting and warmaking, what does a removal of human input mean for the future of warfare as we know it? If autonomous drones, or “killer robots”, cannot identify the difference between combatants and civilians, the potential for disaster seems almost guaranteed (Humble, 2024). The case presented below sheds further light on this ethical dilemma.

Presentation of the Case: AI in the Gaza Strip

As one of the world’s first and foremost users of drone technology in warfare, with widespread utilization dating back to the earliest drone technology in the first Lebanon War in 1982, Israel has also been at the forefront of AWS (The Economist, 2023). In the months (and now years) since the October 7th, 2023 attacks on Israeli targets by Hamas militants, Israel’s invasion into Gaza has been intense and bloody. Made unique by Gaza’s tight, urban environment, the use of drones to construct 3D models of buildings and terrain in order to navigate and locate targets has been fairly well-documented (The Economist, 2023); what has remained less clear, however, is precisely how Israel has gone about determining their targets in Gaza. According to several sources from within Israeli intelligence, however, Israel has used an AI system called Lavender to identify and confirm their targets—over 37,000 of them, in fact, who were identified

based on claims of their individual ties to Hamas (McKernan and Davies, 2024). Lavender is not the only AI tool Israel uses in Gaza: according to Human Rights Watch, Israel uses four different forms of AI in its ongoing assault. These include an unnamed system that tracks mobile phones in order to monitor the evacuation of Palestinians from northern Gaza, a system called “The Gospel”, which generates lists of buildings/structural targets to be hit, the aforementioned “Lavender” system which determines individual military targets, and “Where’s Daddy”, a system that determines when a target is present at a specific location so they can be eliminated (Human Rights Watch 2024).

While each of these systems have drawn attention and criticism from international human rights watchdogs (Human Rights Watch 2024, Davies et. al, 2023), particularly as to whether cell phone data triangulation is accurate enough to inform military decisions (in the case of The Gospel), Lavender has perhaps been the subject of the most outrage and ethical questioning in recent months. The IDF has very generally acknowledged the existence of the system, which it calls a “database whose purpose is to cross-reference intelligence sources” (Human Rights Watch, 2024). Lavender works through machine learning, which assigns Gazans a numerical score measuring the likelihood of their connection to Hamas or another militant group. The system’s algorithm relies on a type of pattern of life analysis, utilizing surveillance data to identify common characteristics in known militants and search for those same characteristics among unknown individuals. The precise characteristics that Lavender searches for are unknown; but given what is currently known about the flawed nature of AI-directed pattern of life analysis and positive unlabeled data, the Lavender system puts civilians at risk (McKernan and Davies 2024; Human Rights Watch 2024).

Similarly to how the guilt of populations is assumed with biometric tracking and predictive policing systems, the constant surveillance of civilians in Gaza pre-assumes their status as members of a militant group. Not only does this run contrary to international humanitarian law, but it opens up the serious possibility of false positives, which can lead to the targeting and deaths of thousands of unaffiliated, unarmed civilians (Human Rights Watch 2024). The broad definitions of terrorism under Israeli law lend no further confidence in the Lavender system, as expansive qualifications for terrorist groups include Palestinian human rights groups, NGOs, and others entirely unaffiliated with terrorist violence; as Human Rights Watch notes, “if similar broad definitions were used in the training of machine learning tools like Lavender, intended to

inform targeting decisions, the outputs would face similar biases and could increase the possibility of civilians being targeted for attack” (2024).

With the extensive use of AI and AWS by the Israeli military to identify targets in Gaza (as well as in the West Bank, which remains subject to Israeli occupation, illegal settlements, and demolition of Palestinian villages (Human Rights Watch, 2024)), serious questions remain as to the ethics of using AI in conflict scenarios—in this case, it seems, with absolute impunity. The IDF’s use of AWS in Gaza is not simply a question of a machine learning system making mistakes and accidentally targeting civilians, as sources speaking to *The Guardian* noted that “We were not interested in killing [Hamas] operatives only when they were in a military building or engaged in a military activity... it’s much easier to bomb a family’s home. The system is built to look for them in these situations” (McKernan and Davies, 2024). AI is not only at risk of unwittingly taking civilian lives in Gaza—evidence suggests the system is specifically designed to take out targets at their most vulnerable, specifically when known and untargeted civilians are highly vulnerable to casualty. With the death toll in Gaza at over 50,000 since October 7th, 2023 (Gaza Ministry of Health quoted in Al Jazeera, 2025), the use of AI in determining further targets—who are often struck with automated drones, such as those used in Ukraine—remains a twisted example of the very harm that runs antithetical to AI’s supposed “harmless” objectives.

Focus Questions

1. What is an AWS, and what are some common examples of these systems?
2. What are some advantages of AWS, and what are some flaws/concerns that have been raised about their use?
3. How has AWS been used in the war in Ukraine?
4. How has the IDF utilized AI in Gaza, and how do these systems work? What data do they use to identify targets?
5. What concerns have been specifically raised about the use of AI in Gaza?

Thematic Reflection and Discussion

International Law and AI in Conflict

As noted in the above background and case study, the relationship between AI in warfare/AWS and international humanitarian law is complicated at best. Though the use of drones and other AI-guided systems are generally subject to the same parameters of international law as other

weapons of war—“namely, the rules of distinction, proportionality, and prohibition of indiscriminate attacks”—there remains a lack of any specific international regulations or guideposts for the use of AI or AWS. With these kinds of weapons systems being used across the globe in numerous different theaters and contexts, law scholars have indicated that a lack of these frameworks and methods for determining ethical use of AI in warfare presents a substantial challenge when attempting to protect international humanitarian law and, by extension, civilian lives in armed conflict.

Discussion Questions

1. Why might an international body, such as the UN, need to construct and pass guidelines surrounding the use of AI and AWS?
2. What do you think the biggest roadblocks to these kinds of regulations are?
3. How might states be regulated and/or held accountable for their use of AI and AWS? What do you think should be permissible under international law?
4. Should private companies be allowed to provide AI systems to states and militaries to be used during war, and under what parameters?
5. How do you think greater transparency could/should be established surrounding the use of AI/AWS?

Pattern of Life Analysis and Flaws in AWS

Perhaps the largest source of concern for those connected to the use of AI in conflict has been related to the inherent flaws and vulnerabilities of AI. Machine learning is imperfect, and numerous issues have been identified related to the ability of AI and/or AWS to accurately identify targets, with or without human input or approval. Systems are vulnerable to poisoning, to inaccurate information, to bias in their algorithms, and often simply lack the human nuance and recognition that complicated warfare scenarios require. Though AI can potentially be used to make warfare less human-intensive and more efficient in identifying targets, is this ethical, or even useful, if the systems are wrong? Further, what are the ethics of using pattern of life analysis in order to identify targets more generally, particularly if the information that systems are being fed is inaccurate or susceptible to bias?

Discussion Questions

1. Do you think the flaws in AI targeting systems mean they should not be used, or are there contexts within which these systems could be advantageous?

2. What are the biggest ethical issues you see with pattern of life analysis?
3. Is there an ethical difference between AI systems that utilize human input and AWS that are entirely autonomous?
4. Do you think machines should be able to identify human beings that should be killed by militaries?

Targeting of Combatants/Civilians in Gaza

The use of AI systems in Israel's invasion of the Gaza Strip has raised questions beyond the pattern of life analysis/AI systems being used to target enemy combatants. In the tight urban setting of the Gaza Strip, Israel uses multiple AI systems to identify targets and track the whereabouts of civilians. Individuals are constantly surveilled, and citizens are consistently and continuously identified as military targets by their participation in unknown behaviors and/or online activity. Sources from within Israeli intelligence have also noted that the AI systems are designed to target individuals, whether they be combatant or civilian, when they are in civilian settings, such as family homes. This is cause for substantial alarm, particularly given that there are no major regulatory frameworks for the use of AI in combat, but also with regard to the international law violations present in the targeting of civilians in combat zones.

Discussion Questions

1. Are there ways to prevent militaries from using AI to surveill and target civilians?
2. Do you think the efficiency of AI can justify "collateral damage" in war?
3. How does the legal structure in Israel influence what individuals are identified as targets?
4. Should the kinds of systems Israel uses in its invasion be used at all during combat?

Bibliography

Al Jazeera. 2025. "Israel's War on Gaza Has Killed 50,000 Palestinians since October 2023." *Al Jazeera*.

<https://www.aljazeera.com/news/2025/3/23/israeli-offensive-in-gaza-has-killed-50000-palestinians-since-october-2023#:~:text=Gaza%E2%80%99s%20Ministry%20of%20Health%20said%20on%20Sunday%20that,the%20Palestinian%20group%20 Hamas%20on%20October%207%2C%202023.>

Davies, Harry, Bethan McKernan, and Dan Sabbagh. 2023. "The Gospel': How Israel Uses AI to Select Bombing Targets in Gaza." *The Guardian*.

[https://www.theguardian.com/world/2023/dec/01/the-gospel-how-israel-uses-ai-to-select-bombing-targets.](https://www.theguardian.com/world/2023/dec/01/the-gospel-how-israel-uses-ai-to-select-bombing-targets)

- Human Rights Watch. 2024. "Questions and Answers: Israeli Military's Use of Digital Tools in Gaza." *HRW*.
<https://www.hrw.org/news/2024/09/10/questions-and-answers-israeli-militarys-use-digital-tools-gaza>.
- Humble, Kristian. 2024. "War, Artificial Intelligence, and the Future of Conflict." *Georgetown Journal of International Affairs*.
<https://gjia.georgetown.edu/2024/07/12/war-artificial-intelligence-and-the-future-of-conflict/>.
- Marr, Bernard. 2024. "How AI Is Used In War Today." *Forbes*.
<https://www.forbes.com/sites/bernardmarr/2024/09/17/how-ai-is-used-in-war-today/>.
- McKernan, Bethan, and Harry Davies. 2024. "'The Machine Did It Coldly': Israel Used AI to Identify 37,000 Hamas Targets." *The Guardian*.
<https://www.theguardian.com/world/2024/apr/03/israel-gaza-ai-database-hamas-airstrikes>.
- Starck, Nick, David Bierbrauer, and Paul Maxwell. 2022. "Artificial Intelligence, Real Risks: Understanding—and Mitigating—Vulnerabilities in the Military Use of AI." *Modern War Institute at West Point*.
<https://mwi.westpoint.edu/artificial-intelligence-real-risks-understanding-and-mitigating-vulnerabilities-in-the-military-use-of-ai/>.
- Stewart, Ruben, and Georgia Hinds. 2023. "Algorithms of War: The Use of Artificial Intelligence in Decision Making in Armed Conflict." *ICRC*.
<https://blogs.icrc.org/law-and-policy/2023/10/24/algorithms-of-war-use-of-artificial-intelligence-decision-making-armed-conflict/>.
- The Economist. 2023. "How Israel Is Using Drones in Gaza." *The Economist*.
<https://www.economist.com/the-economist-explains/2023/12/04/how-israel-is-using-drones-in-gaza>.
- U.S. Department of Defense. 2023. "DoD Announces Update to DoD Directive 3000.09, 'Autonomy In Weapon Systems.'" <https://www.defense.gov/News/Releases/Release/Article/3278076/dod-announces-update-to-dod-directive-300009-autonomy-in-weapon-systems/>.