Secontinuum *Research*

RESEARCH REPORT · SR-2026-01

AI in Recent Conflicts: Ukraine, Gaza, and the Houthi Campaign

How three concurrent wars turned military AI from experiment into infrastructure

SEPTEMBER 2026

PUBLIC · TLP:CLEAR

SECONTINUUM.COM

ABOUT THIS REPORT

Report	SR-2026-01 · version 1.0
Published	20 September 2026
Author	Secontinuum Research
Web edition	www.secontinuum.com/blog/ai-in-recent-conflicts-ukraine-gaza-houthi/
Contact	research@secontinuum.com

HOW THIS REPORT WAS PREPARED

Prepared with AI assistance and reviewed by the Secontinuum team. It draws only on the open sources listed in the references; the authors had no access to classified material.

SUGGESTED CITATION

Secontinuum Research (2026). *AI in Recent Conflicts: Ukraine, Gaza, and the Houthi Campaign*. SR-2026-01. https://www.secontinuum.com/blog/ai-in-recent-conflicts-ukraine-gaza-houthi/

DISCLAIMER

Figures on system performance, targeting accuracy and casualties are taken from open-source and investigative reporting rather than verified official records, and should be treated as contested estimates where the text notes it. This report is published for information only. It is not legal advice and does not represent the position of any government or armed force.

EXECUTIVE SUMMARY

This paper examines the use of artificial intelligence across three concurrent conflicts (the Russia-Ukraine war, the Israel-Gaza war, and the Houthi campaign against Red Sea shipping) as complementary case studies spanning state and non-state actors, open-battlefield, urban and maritime environments, and applications ranging from ISR and autonomous navigation to population-scale target generation and asymmetric terminal guidance. Drawing on open-source and investigative reporting, it finds that resource constraint rather than resource abundance has driven most visible AI adoption; that the stage of the kill chain at which AI operates predicts legal and ethical stakes better than technological sophistication does; and that capability and countermeasure diffusion across actors is outpacing the non-binding governance frameworks meant to constrain it. No confirmed instance of a fully autonomous, human-out-of-the-loop lethal engagement appears in any of the three cases, but human review in at least one case is reported to have compressed to a matter of seconds.

KEY FINDINGS

01 **Constraint, not abundance, drives adoption**

Ukraine automated because it was short of soldiers and materiel; the Houthis because they lack ISR and precision-strike infrastructure. In each case AI closes a gap between what a belligerent wants to do and what its people can do unassisted.

02 **Where AI sits matters more than how advanced it is**

Upstream, population-scale target generation raises the sharpest distinction and proportionality questions. Downstream terminal guidance raises narrower concerns about reliability. Kill-chain position predicts the legal stakes better than sophistication does.

03 **The human is still in the loop, but the loop is shrinking**

No confirmed fully autonomous lethal engagement appears in any of the three conflicts. Reported human review of machine-generated targets has nonetheless compressed to roughly 20 seconds in one case.

04 **Diffusion is outrunning governance**

Components, techniques and countermeasures cross the state/non-state divide in weeks. REAIM and the UN processes remain non-binding, and the 2026 target for autonomous-weapons rules is unlikely to be met.

CONTENTS

1. Introduction 5

2. Background and Context 5

3. Case Study: Russia-Ukraine War 6

4. Case Study: Israel-Gaza Conflict 7

5. Case Study: Houthi Operations in Yemen and the Red Sea 8

6. Cross-Cutting Technology Themes 9

7. Legal and Ethical Implications 11

8. Policy and Governance Implications 12

9. Comparative Analysis and Emerging Patterns 13

10. Conclusion and Future Outlook 14

References 16

1. Introduction

Between 2022 and 2026, artificial intelligence moved from a peripheral, experimental role in warfare to a load-bearing component of three concurrent conflicts: the Russia-Ukraine war, the Israel-Hamas war in Gaza, and the Houthi campaign against shipping and military targets in the Red Sea and Bab al-Mandab Strait. Each conflict has become a distinct proving ground for a different class of military AI application, and together they offer the fullest real-world picture available of how machine learning, computer vision, and autonomous or semi-autonomous systems are reshaping the conduct of war.

This paper defines "AI" pragmatically, as it is used by militaries and reported by researchers in these conflicts: machine learning models and computer vision systems used for target recognition and classification; algorithmic decision-support and target-generation systems that rank or nominate targets for human review; autonomous or semi-autonomous navigation and terminal guidance in uncrewed systems; and AI-assisted tools for signals analysis, sensor fusion, and information operations (including generative AI used in disinformation). It deliberately excludes older forms of automation (such as fire-and-forget guided munitions with pre-AI seekers) except where they form a direct lineage to newer AI-enabled variants.

The three cases were selected because they are analytically complementary rather than redundant. Ukraine represents a high-intensity, resource-constrained state actor rapidly iterating AI-enabled uncrewed systems under direct electronic-warfare pressure. Gaza represents a technologically advanced state military using AI-driven decision-support and target-generation systems at population scale in a dense urban environment. The Houthi campaign represents a non-state, Iran-backed actor achieving asymmetric effects against far better-resourced adversaries using comparatively low-cost, partly AI-guided drones and missiles built from commercial and transferred components. Read together, they span the spectrum from state to non-state actors, from open battlefield to urban and maritime environments, and from AI as a targeting aid to AI as a force multiplier for materially weaker belligerents.

The paper draws on open-source investigative reporting, think-tank and international-organization analysis, and official and leaked military documentation as reported by credible outlets; it does not have access to classified sources, and figures on system performance and casualties should be read as reported estimates rather than verified counts. Section 2 provides brief technological and historical context. Sections 3 through 5 present the three case studies. Section 6 draws out cross-cutting technology themes. Sections 7 and 8 examine legal, ethical, and policy implications. Section 9 offers a comparative analysis, and Section 10 concludes with an assessment of near-term trajectory.

2. Background and Context

Military interest in autonomous and AI-assisted targeting predates the current conflicts by decades, from Cold War-era anti-radiation missiles to the DARPA-funded computer-vision research of the 2000s and 2010s. What changed after roughly 2020 was the convergence of four civilian-sector developments that lowered the cost of fielding AI at scale: cheap, capable commercial drones and components; consumer-grade GPUs and edge-computing chips (such as Nvidia's Jetson line) powerful enough for onboard inference; open-source computer-vision and machine-learning frameworks that shortened development cycles from years to weeks; and cloud infrastructure and commercial satellite imagery that made large-scale sensor fusion and data labeling feasible for militaries without a national AI-industrial base.

These enablers diffused unevenly. Well-resourced state militaries such as Israel's built proprietary, classified AI targeting infrastructure integrated with existing intelligence architecture. Ukraine, lacking that depth of legacy infrastructure but facing acute manpower and materiel constraints, instead built an ecosystem of small, fast-iterating domestic defense-tech firms and volunteer engineering groups producing lightweight, modular AI software compatible with commodity drone hardware [1]. Non-state and state-sponsored actors such as the Houthis took a third path, relying on transferred and reverse-engineered Iranian missile and drone designs supplemented with commercial off-the-shelf navigation, sensor, and communications components [11] [12].

A further contextual factor is the compressed adaptation cycle these conflicts have produced. Where Cold War weapons development ran on multi-year procurement cycles, Ukrainian and Russian forces now iterate drone software and counter-drone electronic warfare on a cycle of weeks, described by analysts as an accelerating arms race between AI-enabled targeting and AI-resistant jamming [2]. This compression is itself a defining feature of AI's role in recent conflict: it is not only what AI does on the battlefield, but how quickly opposing forces can retrain, redeploy, and counter it, that increasingly determines battlefield effect.

3. Case Study: Russia-Ukraine War

Ukraine has positioned removing soldiers from direct combat exposure through AI-enabled uncrewed systems as a central strategic objective, pursued not through a single formalized doctrine but through a distributed ecosystem of domestic defense-tech firms, military units, and volunteer engineering groups [1]. In 2024 alone Ukraine produced roughly 2 million drones, including about 1.5 million first-person-view (FPV) drones, with domestic production accounting for over 96 percent of the military's drone supply; that same year the armed forces began purchasing an initial 10,000 AI-enhanced drones as a first step toward broader autonomous adoption.

ISR and sensor fusion. The Delta command-and-control system integrates tens of terabytes of daily intelligence data from multiple sources. Purpose-built AI tools feed into it: the Zvook acoustic network, covering roughly 20,000 square kilometers with passive sensors costing about $500 each, detects and classifies aerial threats (distinguishing drones, cruise missiles, and helicopters) up to 4.8 km away with a 1.6 percent false-positive rate; the crowdsourced ePPO mobile app, with over 600,000 downloads, routes AI-processed civilian threat reports to air-defense units within 2 to 7 seconds; and the Griselda platform applies AI transcription and entity recognition to intercepted communications and chat data, cutting the associated analyst workload by an estimated 99 percent.

Automatic target recognition and terminal guidance. Systems such as the compact ZIR autonomy module and Skydio X10D drones (roughly 1,000 deployed) use onboard computer vision, including models adapted from commercial sports-tracking software, to maintain a lock on a target without a continuous radio or GPS link, a direct response to Russian jamming (see below). Ukrainian firm Fourth Law's roughly $70 AI vision modules reportedly raised one brigade's FPV strike success rate from about 20 percent to 80 percent, and AI-enabled autonomous terminal guidance more broadly has been reported to raise engagement success from a 10-20 percent baseline to 70-80 percent, reducing the number of drones needed to strike a target from 8-9 to 1-2.

> *From 8 or 9 drones per target down to 1 or 2.*

Electronic warfare arms race. Russian systems including the Krasukha-4 broadband jammer, R-330Zh Zhitel VHF/UHF jammer, and the GPS-focused Pole-21 degraded Excalibur precision-munition accuracy by up to 90 percent in heavily contested sectors and cut FPV drone hit rates from 40-60 percent in 2023 to 20-30 percent or lower by mid-2024. Ukraine's response combined AI-based visual-inertial navigation (allowing drones to navigate without GPS) with a shift to fiber-optic-guided FPV drones immune to radio jamming; by late 2025 Russia itself was reportedly producing an estimated 50,000 or more fiber-optic FPV drones per month, underscoring how quickly countermeasures diffuse to both sides [2].

Autonomy in practice. Despite this, both Ukrainian and independent assessments describe current systems as exhibiting "partial," not full, autonomy: AI enhances navigation, tracking, and last-mile guidance, but human operators retain engagement authority and can override automated functions. The December 2024 Lyptsi operation, involving dozens of uncrewed ground vehicles and FPV drones in a fully unmanned multi-domain assault with no infantry participation, is frequently cited as a landmark, but the vehicles involved remained manually operated rather than autonomously targeting. Ukrainian officials describe extending autonomy to coordinated multi-drone swarming and to ground, sea, and undersea platforms as the next frontier, not yet realized at scale.

4. Case Study: Israel-Gaza Conflict

Where the Ukrainian case is defined by distributed, fast-iterating tactical AI, the Israeli case is defined by centralized, intelligence-integrated AI used for target generation at population scale. Reporting based on interviews with Israeli intelligence officers, principally by +972 Magazine and Local Call and corroborated by subsequent investigations, describes an "AI factory" approach built around three linked systems [5] [6] [7]:

- **Lavender** is a machine-learning system that draws on mass surveillance data covering most of Gaza's roughly 2.3 million residents to rate individuals from 1 to 100 on likelihood of affiliation with Hamas or Palestinian Islamic Jihad military wings, effectively generating a ranked kill list.

- **The Gospel (Habsora)** performs an analogous function for physical structures and infrastructure rather than individuals, nominating buildings for strike.
- **"Where's Daddy?"** is an automated tracking tool that alerts operators when a Lavender-marked individual enters a residence, enabling strikes timed to when the target and, frequently, family members are present.

According to this reporting, Lavender marked approximately 37,000 Palestinians as potential targets during the early weeks of the war, with sweeping authorization to act on those marks granted roughly two weeks in after a sample audit found about 90 percent accuracy against a human-defined ground truth, meaning an estimated 1 in 10 marked individuals were not, in fact, militants. Reported sources of error included flagging people who merely shared communication patterns or previously owned phone numbers associated with actual operatives, and misidentifying police and civil-defense personnel. Human review of each machine-generated target was reported to average roughly 20 seconds, often limited to confirming the target's gender.

> *Human review of each machine-generated target was reported to average roughly 20 seconds.*

The reporting further describes permissive, pre-authorized collateral-damage thresholds tied to a target's assigned rank: up to 15-20 civilian deaths were reportedly authorized per strike on a junior operative, and over 100 for senior commanders, thresholds described by sources as exceeding those used in prior U.S. operations against ISIS. Because collateral-damage estimates were calculated using outdated, pre-war population and evacuation figures rather than real-time verification, and because strikes on "Where's Daddy"-flagged homes proceeded without confirming the target was actually present, the investigation attributes a significant share of Gaza civilian casualties in the war's first six weeks (reported as approximately 15,000 deaths, nearly half the total for that period) to this targeting pipeline, though these figures come from single-source reporting and should be treated as contested estimates rather than verified counts.

Israeli military officials have disputed aspects of this characterization, describing Lavender as one input among several subject to independent verification rather than an autonomous targeting system, and emphasizing that final strike decisions remain with human commanders. Independent legal scholars nonetheless note that even accepting the military's characterization, the volume of AI-generated nominations, the speed of review, and the pre-set casualty thresholds raise distinct questions under the law of armed conflict's distinction and proportionality requirements that are examined further in Section 7 [7].

5. Case Study: Houthi Operations in Yemen and the Red Sea

Since November 2023 the Houthi movement has conducted over 100 attacks on commercial and naval shipping transiting the Red Sea and Bab al-Mandab Strait, cutting Suez Canal crossings from roughly 2,068 in November 2023 to about 877 by October 2024 and prompting the multinational Operation Prosperity Guardian naval response [9] [10]. This is the paper's clearest example of a materially weak, non-state actor generating strategic-level disruption through a layered, partly AI- and automation-enabled arsenal rather than any single breakthrough capability.

Arsenal and guidance technology. Open-source military analysis identifies at least six distinct Houthi anti-ship ballistic missile designs, including the Asef (an Iranian Fateh-313 derivative with an electro-optical/infrared seeker and roughly 280-mile range) and the Tankil (based on Iran's Raad-500, ranging roughly 310 miles), alongside multiple anti-ship cruise missiles derived from Iranian Noor/Ghader and Quds Z-0 designs, some with radar-homing or electro-optical/infrared terminal seekers and land-attack/anti-ship dual capability [11]. As of January 2024 the Houthis had become the first force in history to fire anti-ship ballistic missiles in combat. Their one-way-attack drones, including the "Yafa" model used in a July 2024 strike on Tel Aviv, and a range of Iranian Shahed-derived platforms, rely on a mix of inertial, satellite, and in some variants electro-optical terminal guidance; UN and industry investigations have repeatedly documented Iranian-sourced and reverse-engineered commercial components, including commercial-grade navigation chips, engines, and electronics, inside interdicted weapons shipments [12].

Targeting without exquisite ISR. Unlike the Ukrainian and Israeli cases, the Houthis lack dedicated satellites or maritime patrol aircraft. They compensate with a lower-technology but effective sensing layer: small UAVs, civilian-flagged scout vessels, commercial ship-tracking (AIS) data, and intelligence support from the Iranian vessel *Behshad*, widely assessed to function as an Islamic Revolutionary Guard Corps forward intelligence and targeting platform in the Red Sea [11]. This illustrates a distinct pattern from the other two cases: rather than AI-driven sensing feeding AI-driven targeting end to end, the Houthi model pairs largely human- and commercially-sourced targeting data with automated or semi-automated terminal guidance, achieving asymmetric effect through volume, persistence, and the difficulty and expense of intercepting low-cost threats with high-cost naval air-defense systems.

Escalation and countermeasures. Attacks evolved across expanding phases from late 2023 to 2025, broadening from Israel-linked vessels to any ship with U.S. or British ties or Israeli port calls, and extending to direct strikes on Israeli territory, including a 2025 drone attack reaching Tel Aviv's airport vicinity and an attack on Eilat. Despite roughly 62 missiles and drones fired across 26-plus incidents by January 2024 alone, no vessel had been confirmed sunk in that period, reflecting both the limits of the Houthis' targeting precision and the effectiveness of coalition naval defenses, though the economic and shipping-diversion costs have been substantial regardless of kinetic success rate.

6. Cross-Cutting Technology Themes

Levels of autonomy. None of the three cases involves weapons that autonomously select and engage targets without a human decision point built in somewhere in the kill chain, but the location of that decision point varies sharply. Ukraine keeps a human in the loop at engagement while automating navigation and tracking; Israel's reported systems keep nominal human review at target generation while compressing that review to seconds; the Houthis apply automation mainly to terminal guidance rather than target selection, which remains a separate, slower, human- and sensor-driven process. This spread illustrates that "autonomy" in current conflict is rarely all-or-nothing but is instead distributed unevenly across the sensing, deciding, and engaging stages of the kill chain.

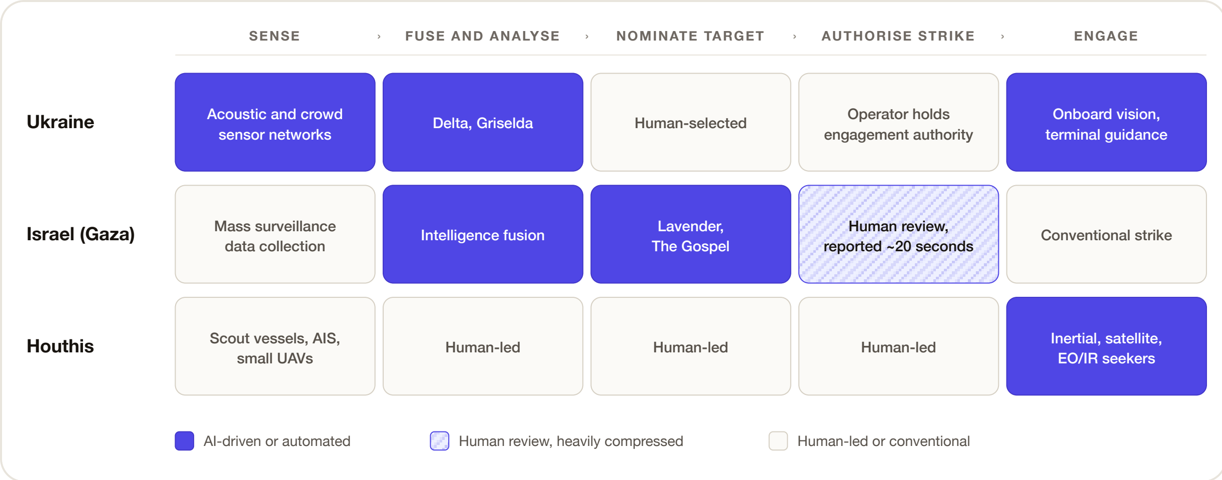

Figure 1. Where AI sits in the kill chain for each belligerent, as described in open-source reporting. Source: Secontinuum Research.

Computer vision and target recognition. All three theaters show computer-vision models, several adapted from civilian applications (sports analytics, consumer robotics), performing classification tasks under degraded conditions: GPS-denied, jammed, or low-bandwidth environments in Ukraine; large-scale pattern-of-life classification against surveillance data in Gaza; and, more provisionally, electro-optical/infrared terminal seekers in Houthi missiles. The common technical thread is a shift from GPS- or radio-dependent guidance toward onboard, model-based perception that is harder to jam but also harder to audit or explain after the fact.

Sensor fusion and ISR. Ukraine's Delta system and Griselda platform, and Israel's integration of surveillance data into Lavender and the Gospel, represent the most mature examples of AI used to fuse and triage intelligence volumes beyond human analytic capacity. The Houthi case is the outlier: lacking comparable ISR infrastructure, the group substitutes a lower-technology sensing layer (scout vessels, AIS tracking, Iranian naval support), suggesting that AI-enabled sensor fusion remains a state-level capability gap not yet closed by widely available commercial technology.

Electronic warfare and counter-UAS. The Ukraine theater has produced the most developed AI-versus-EW dynamic of the three, with jamming and AI-based counter-jamming (visual-inertial navigation, fiber-optic guidance, mesh networking) iterating on a cycle of weeks. Coalition naval forces face an analogous but less AI-centric problem intercepting Houthi threats, relying more on kinetic and electronic point-defense than on AI-driven countermeasures to date.

AI-enabled cyber operations. Cyber operations accompany all three conflicts (disruptive attacks on Ukrainian infrastructure, hacktivist activity tied to the Gaza war, and Houthi-linked or Iran-aligned network intrusions), though open-source reporting on AI-specific enhancement of offensive cyber tooling (as opposed to AI-assisted defense and anomaly detection) remains comparatively thin next to the well-documented kinetic AI applications above; this is an area where classification gaps likely understate real usage.

Information warfare and generative AI. Both the Ukraine and Gaza conflicts have produced well-documented instances of AI-generated or AI-manipulated imagery deployed for propaganda and disinformation, including fabricated images of civilian casualties (notably infants) circulated to shape international opinion, and synthetic audio and video used in both directions of the Ukraine-Russia information contest [17] [18]. Researchers describe this less as a decisive new weapon than as an accelerant that lowers the cost and increases the volume and believability of information operations that predate generative AI, while also complicating fact-finding by making authentic atrocity evidence easier to dismiss as fabricated, a dynamic sometimes called the "liar's dividend."

7. Legal and Ethical Implications

The International Committee of the Red Cross, in its most recent position paper on autonomous weapon systems (AWS), frames the central legal difficulty as a mismatch between what international humanitarian law (IHL) requires and what current AI systems can reliably do [13]. The principle of distinction requires context-specific human judgment to separate combatants from civilians and from persons hors de combat, a judgment the ICRC argues AWS cannot reliably perform, particularly for objects that are military objectives "by use" rather than by nature. Proportionality assessments, which weigh anticipated civilian harm against concrete military advantage, become especially difficult when a system is activated without knowing in advance exactly which targets it will engage, a problem magnified in dense urban environments such as Gaza. Because AWS select or nominate targets after activation, responsibility for a given strike's lawfulness becomes harder to trace to an identifiable human decision, producing what the ICRC and other bodies describe as an accountability gap.

The ICRC's proposed response is two absolute prohibitions (on AWS whose behavior cannot be sufficiently understood or predicted, and on any system designed to target humans directly) plus strict operational limits (restricted target types, duration, geographic scope, and mandatory human supervision with real-time intervention capability) on systems that fall short of those prohibitions. It draws an explicit historical parallel to anti-personnel landmines, arguing that early, pre-emptive prohibition proved more achievable than restriction after widespread fielding, as with the 1997 Ottawa Treaty.

Applied to the case studies, the legal picture differs by case. The Ukrainian systems described in Section 3 retain human engagement authority and are, on available reporting, closer to the "permissible with constraints" category the ICRC describes, though questions remain about how meaningfully operators can review AI-flagged targets under the time pressure of contested electromagnetic environments. The Gaza case, as detailed in Section 4, sits closer to the center of the ICRC's concerns: reported 20-second human reviews of machine-generated nominations, pre-authorized civilian-casualty thresholds set before specific strikes were contemplated, and population estimates disconnected from real-time conditions collectively describe a targeting process where the human role, while formally retained, may not constitute the substantive judgment IHL's distinction and proportionality rules presuppose, a concern raised directly by legal scholars reviewing the reporting [7] [8]. The Houthi case raises a related but distinct question: because Houthi missiles and drones frequently miss their intended military or dual-use targets and have struck civilian and third-country shipping, the concern is less about AI displacing human judgment and more about the combination of automated terminal guidance and comparatively unsophisticated targeting data producing indiscriminate effect as a practical matter, regardless of the attacker's formal targeting process or its adherence (or lack thereof) to IHL as a non-state actor not party to the relevant treaties in the same way as states.

Across all three cases, a common ethical concern is what several scholars term "automation bias": the tendency of human reviewers, even when nominally in the loop, to defer to a machine's output rather than exercise independent judgment, especially under time pressure or high operational tempo. The Gaza reporting describes this dynamic explicitly, and Ukrainian and Houthi-facing analysts increasingly flag it as a risk as review timelines compress across all forms of AI-assisted targeting.

8. Policy and Governance Implications

International governance of military AI has developed largely through non-binding, multi-stakeholder processes rather than treaty law. The Summit on Responsible AI in the Military Domain (REAIM) is the most prominent such process, convening in The Hague (2023), Seoul (2024), and A Coruña, Spain (February 2026). Its 2024 "Blueprint for Action" was endorsed by more than 60 countries but notably not by China, which objected to language requiring human involvement in nuclear-weapons decisions; Russia has not participated, and U.S. engagement has diminished under the current administration's emphasis on what it describes as "hard-nosed realism" over multilateral AI constraints [16]. A parallel track runs through the UN Group of Governmental Experts on Lethal Autonomous Weapons Systems under the Convention on Certain Conventional Weapons, and through UN General Assembly resolutions; in May 2025 the UN Secretary-General renewed a call for a legally binding instrument on autonomous weapons, and the ICRC and UN Secretary-General have jointly urged states to conclude negotiations on AWS-specific rules by the end of 2026, a deadline that, as of this writing, appears unlikely to be met given continuing disagreement over verification, scope, and definitions [14] [15].

The three case studies illustrate why consensus has proven difficult. Ukraine's experience has made rapid, decentralized AI adoption a matter of national survival, giving Kyiv and its Western partners limited appetite for restrictions that would slow battlefield iteration; this has fed a broader Western policy shift, visible in loosened U.S. and European export-control postures for military drone and AI components, toward prioritizing speed of fielding over precaution. Israel's use of AI targeting systems, conducted by a state with an advanced indigenous defense-AI industry operating largely outside public or treaty-based oversight, illustrates the limits of voluntary governance frameworks when a state judges a capability to confer decisive advantage. The Houthi case illustrates a different governance failure: the group's arsenal depends on the transfer and reverse-engineering of components subject to existing UN and unilateral sanctions and arms embargoes on Iran and Yemen, indicating that export-control and non-proliferation regimes, not AI-specific governance, remain the more relevant and already-strained lever for a non-state actor case, and that their persistent porousness is a policy failure distinct from the AWS-specific debate.

Taken together, the policy landscape by late 2026 is characterized by a growing normative consensus, articulated through REAIM, ICRC advocacy, and UN processes, that unconstrained autonomous targeting is undesirable, running against a practical reality in which the states and non-state actors actually fighting these three wars face strong incentives to adopt AI capabilities faster than any binding international instrument is likely to constrain them, and in which existing arms-control tools (export controls, sanctions, interdiction) remain more operative than AI-specific rules for at least the non-state-actor case.

9. Comparative Analysis and Emerging Patterns

Three patterns emerge from reading the cases side by side rather than in isolation. Table 1 summarizes the comparison.

	UKRAINE	ISRAEL (GAZA)	HOUTHIS
Actor	Resource-constrained state	Technologically advanced state	Non-state, Iran-backed
Environment	Open battlefield under heavy EW	Dense urban	Maritime chokepoint
Where AI is concentrated	Navigation and terminal guidance, plus ISR fusion	Target generation and intelligence fusion	Terminal guidance only
Human decision point	Operator holds engagement authority	Nominal review of nominations, reported at about 20 seconds	Targeting is human-led and separate from guidance
What AI substitutes for	Manpower and materiel	Analyst time against intelligence volume	ISR and precision-strike infrastructure
Principal legal concern	Quality of review under time pressure	Distinction and proportionality at scale	Indiscriminate effect from poor targeting data

Table 1. The three cases compared. Source: Secontinuum Research, from the reporting cited in Sections 3 to 5.

First, resource constraint, not resource abundance, has been the strongest driver of visible AI adoption. Ukraine's innovation was forced by manpower shortages and materiel scarcity relative to Russia; the Houthis' reliance on automation in terminal guidance substitutes for the ISR and precision-strike infrastructure they cannot otherwise field. Israel is the partial exception, applying AI not to overcome resource scarcity but to compress the time and labor cost of processing an intelligence volume that would otherwise outstrip human analytic capacity, suggesting that AI adoption in this period has been driven less by any single motive than by mismatches, of different kinds, between what a belligerent wants to do and what its available human resources can do unassisted.

Second, the level at which AI operates in the kill chain tracks the type of actor more than the sophistication of the technology available to it. Israel, a technologically advanced state, concentrates AI at the target-generation and intelligence-fusion stage, upstream of the strike decision. Ukraine concentrates it at the navigation and terminal-guidance stage, downstream of a human-selected target. The Houthis concentrate it almost entirely at terminal guidance, with targeting itself remaining a comparatively low-tech, human- and open-source-data-driven process. This suggests that where AI is inserted in the targeting process is a better predictor of the resulting legal and ethical stakes than how advanced the AI itself is: upstream, population-scale target generation raises the sharpest distinction and proportionality concerns (Section 7), while downstream terminal-guidance automation raises comparatively narrower, if still real, concerns about reliability and discrimination in cluttered or contested environments.

Third, all three conflicts show technology and countermeasure diffusing across the state/non-state divide faster than governance frameworks can track. Commercial drone and AI components flow, often through gray or black markets, from a global civilian technology base into Houthi hands via Iranian transfer networks; Ukrainian open-source AI software techniques are shared publicly and studied by researchers and, in adapted form, by other conflicts' belligerents; and Israeli AI-targeting practices are cited by other militaries and by arms-control advocates alike as a template to be emulated or foreclosed. The result is that no single case can be treated as sui generis: capabilities and practices demonstrated in one theater are already shaping expectations, doctrine, and procurement in the others, a diffusion dynamic that outpaces the REAIM and UN governance processes described in Section 8 by a wide and, on current trends, widening margin.

10. Conclusion and Future Outlook

By late 2026, AI in the Russia-Ukraine war, the Gaza conflict, and the Houthi campaign has moved well past the experimental phase into routine, load-bearing use across ISR, target generation, terminal guidance, and information operations, without producing anything resembling a fully autonomous, human-out-of-the-loop weapon in confirmed combat use in any of the three cases. The more consequential change has been the erosion of the time and resource cost of fielding militarily significant AI capability, which has let a resource-constrained state (Ukraine), a technologically advanced state (Israel), and a non-state armed group (the Houthis) each build materially different but individually potent AI-enabled capabilities within roughly the same multi-year window.

Three trajectories bear watching. Governance is likely to remain a patchwork of non-binding norms (REAIM), slow-moving UN treaty processes, and case-specific export-control enforcement rather than a single binding instrument, at least through the ICRC and UN Secretary-General's 2026 target date and probably beyond it. The autonomy frontier, particularly coordinated drone swarming and further compression of human review at the target-generation stage, is the area both Ukrainian officials and independent analysts identify as the next likely point of escalation, and the one where the legal and ethical concerns raised in Section 7 will be tested most directly. And diffusion will very likely continue to outpace regulation: techniques and components demonstrated or transferred in any one of these three conflicts should be expected to appear, adapted, in the next one, making the comparative reading offered in Section 9 as much a forecasting exercise as a retrospective one.

References

All links accessed 19 September 2026.

1. Kateryna Bondar, "Ukraine's Future Vision and Current Capabilities for Waging AI-Enabled Autonomous Warfare," Center for Strategic and International Studies, 6 March 2025.
2. Clara Kaluderovic, "The Coming Compute War in Ukraine," Atlantic Council, 16 March 2026.
3. Federica Marsi and Mansur Mirovalev, "'Attacks will be fully autonomous': Russia, Ukraine race towards AI warfare," Al Jazeera, 14 September 2026.
4. Alex Carter, "How Russia's Electronic Warfare Blinded Ukrainian Drones, and How Ukraine Fought Back," Military Machine, 3 April 2026.
5. Yuval Abraham, "'Lavender': The AI machine directing Israel's bombing spree in Gaza," +972 Magazine with Local Call, 3 April 2024.
6. Elizabeth Dwoskin, "Israel built an 'AI factory' for war. It unleashed it in Gaza.," The Washington Post, 29 December 2024.
7. Michael N. Schmitt, "The Gospel, Lavender, and the Law of Armed Conflict," Lieber Institute for Law and Warfare, West Point, 28 June 2024.
8. Atay Kozlovski, "When Algorithms Decide Who is a Target: IDF's use of AI in Gaza," Tech Policy Press, 13 May 2024.
9. Noam Raydan and Farzin Nadimi, "Houthi Shipping Attacks: Patterns and Expectations for 2025," The Washington Institute for Near East Policy, 16 December 2024.
10. Council on Foreign Relations, "Conflict in Yemen and the Red Sea," Global Conflict Tracker, updated 3 August 2026.
11. Joseph Trevithick, "The Anti-Ship Missile Arsenal Houthis Are Firing Into The Red Sea," The War Zone, 10 January 2024.
12. Nada AlTaher, "Houthis expand arsenal with drones and missiles from Iran," The National, 11 May 2026.
13. International Committee of the Red Cross, "Autonomous Weapon Systems and International Humanitarian Law: Selected Issues," position paper, December 2025.
14. UN News, "'Politically unacceptable, morally repugnant': UN chief calls for global ban on 'killer robots'," United Nations, 14 May 2025.
15. Alexander Kmentt, "Geopolitics and the Regulation of Autonomous Weapons Systems," Arms Control Today, January/February 2025.
16. Tobias Vestner and Simon Cleobury, "The Third Summit on Responsible AI in the Military Domain (REAIM)," Just Security, 2 February 2026.
17. Gretel Kahn, "Trolling, memes and deepfakes: How AI is thickening the fog of war," Reuters Institute for the Study of Journalism, 4 May 2026.
18. David Klepper, "Fake babies, real horror: Deepfakes from the Gaza war increase fears about AI's power to mislead," Associated Press via WHYY, 28 November 2023.

Machine-Speed SOC for the Agentic Era.

Secontinuum builds Vynnn, the Autonomous SOC: AI agents that detect, investigate and contain threats end to end, with human approval for high-impact actions.

Secontinuum Research studies how autonomous systems change the balance between attackers and defenders, on the battlefield and in the enterprise.

RESEARCH@SECONTINUUM.COM