AI Decision-Support Systems in Warfare: Lessons from the Gaza Case Study

by Emmanuelle Low

King's Policy Journal

KCL Policy Research Centre

Centre for Technology and Artificial Intelligence

Word Count: 2315

January 2026

AI Decision-Support Systems in Warfare: Lessons from the Gaza Case Study by Emmanuelle Low

Abstract

Artificial intelligence has become central to debates about contemporary warfare, humanitarian law and accountability. This paper examines the use of AI-enabled decision-support systems (AI-DSS) in military targeting through a focused case study of their reported deployment by the Israel Defense Forces (IDF) in Gaza. While similar systems have been integrated into other conflicts, the Gaza campaign offers a particularly salient case for analysing how AI-DSS functions when embedded across a digitised targeting cycle operating at scale.

Investigative reporting by +972 Magazine and Local Call, based on interviews with IDF intelligence officers, describes systems such as Lavender, Gospel, and "Where's Daddy?" as generating large volumes of targets under conditions of limited human review and permissive civilian-casualty thresholds. The IDF characterises these tools as intelligence-management systems designed to assist analysts and improve precision, though independent verification of their operational role remains limited.

Rather than adjudicating contested facts about individual strikes, this paper analyses what the reported use of AI-DSS in Gaza reveals about structural risks in AI-assisted targeting. It argues that, when combined with organisational incentives to accelerate operations and permissive rules of engagement, AI-DSS tend to expand and compress the targeting process rather than refine it. This dynamic raises systematic challenges for distinction, proportionality, and accountability under international humanitarian law, with implications extending beyond the Gaza context to the future regulation of military AI.

AI Decision-Support Systems

AI-enabled decision-support systems (AI-DSS) are increasingly embedded across contemporary military command-and-control and targeting architectures. Klonowska defines AI-DSS as systems that apply machine-learning or statistical techniques to large datasets in order to generate probabilistic assessments, correlations or ranked recommendations that assist human decision-makers across the chain of command (Klonowska, 2023). Crucially, AI-DSS differ from autonomous weapons systems in that they do not independently select or engage targets; instead, they structure how information is processed, prioritised and presented to human operators.

Within the targeting cycle, AI-DSS typically perform three core functions. First, they aggregate and fuse heterogeneous intelligence streams, including drone imagery, intercepted communications,

metadata, and social-network analysis (Gusterson, 2016). Second, they identify behavioural or relational patterns associated with suspected militants or military objectives, often through probabilistic classification rather than definitive identification (Moyn, 2021). Third, they support downstream decisions by prioritising targets, estimating collateral damage, or recommending munitions, thereby shaping the tempo and scope of lethal operations (Allen and Husain, 2017). While human authorisation formally remains in place, these systems significantly influence the epistemic environment within which lethal decisions are made (Cummings, 2023).

AI-DSS are not inherently unlawful or inhumane. In principle, they could reduce error, strengthen verification, and enhance compliance with international humanitarian law by improving situational awareness and cross-checking intelligence (ICRC, 2019). Whether they do so, however, depends less on technical capability than on institutional design choices: how models are trained and validated, how uncertainty is communicated to operators, how human review is structured, and what operational incentives govern their use (Roff and Moyes, 2016). Historically, AI-DSS extend earlier forms of algorithmic warfare, particularly post-9/11 metadata-driven targeting and signature strikes, which similarly relied on probabilistic inference rather than individualised knowledge (Gusterson, 2016). As subsequent sections demonstrate, the risks associated with AI-DSS emerge most sharply when such systems are embedded within organisational contexts that prioritise speed, scale, and target generation over deliberation and restraint (Sylvia, 2024).

Case Study: AI-DSS in the Israel–Palestine War

Investigative reporting by +972 Magazine and Local Call provides the most detailed public account to date of the reported use of AI-enabled decision-support systems by the Israel Defense Forces (IDF) during the Gaza campaign. Based on interviews with six current and former IDF intelligence officers, these reports describe a set of systems - commonly referred to as Lavender, Gospel, and “Where’s Daddy?”- as operating together as an integrated targeting pipeline (Abraham, 2024). While the operational details of these systems cannot be independently verified, the reporting offers a valuable window into how AI-DSS may function when embedded at scale within a digitised kill chain.

According to these accounts, Lavender generates probabilistic assessments of individuals based on metadata, communication patterns and social-network associations, producing large volumes of suspected operatives rather than definitive identifications. “Where’s Daddy?” is described as a tracking tool that identifies when flagged individuals are present at specific locations, including private homes, while Gospel and related systems such as Fire Factory automate the nomination of building targets and the generation of

strike plans. Taken together, these systems are reported to structure the flow of targeting decisions by prioritising who is flagged, when they are locatable and which physical sites are nominated for attack.

Additional technologies reportedly deployed by the IDF - including the Red Wolf and Blue Wolf facial-recognition systems, the semi-autonomous Jaguar ground robot, and Sentry Tech remote-controlled gun towers - illustrate the broader surveillance and targeting ecosystem within which these AI-DSS operate. The analytical significance of these tools lies not in their autonomy, but in their cumulative effect: the fusion of multiple data streams into a continuous, machine-mediated process of identification, tracking, and strike nomination.

The most consequential feature highlighted by this reporting is the transformation of the pace and scale of targeting. +972 reports that Lavender flagged up to 37,000 individuals as potential targets, a figure that exceeds pre-war public estimates of Hamas's organised fighting force, commonly placed at 25,000–30,000 (McKernan and Davies, 2024). This expansion aligns with long-standing institutional ambitions within the IDF to accelerate target production. Former Chief of Staff Aviv Kochavi stated in 2021 that the military sought to generate "as many targets in a month as it once generated in a year," signalling an organisational commitment to speed and volume rather than case-by-case deliberation (Kochavi, 2021).

Crucially, investigative sources describe human review as increasingly compressed. Reported approval processes lasting seconds rather than minutes suggest a shift from substantive verification toward procedural confirmation, raising the risk that human oversight becomes formal rather than meaningful (Abraham, 2024). Whether or not such practices were universal, the reported pattern illustrates a structural vulnerability of AI-DSS: when systems generate targets at scale, human review is pressured to adapt to machine tempo rather than constrain it.

The humanitarian consequences of this mode of operation are difficult to disentangle from broader strategic and political decisions. Nonetheless, reported practices - including tolerated civilian-casualty thresholds of 15–20 civilians for lower-level operatives, higher thresholds for senior targets, and extensive use of unguided munitions - suggest an operational environment in which civilian harm was anticipated rather than exceptional (Abraham, 2024; CNN, 2023). UN OCHA reports that during the first month of the campaign, over 6,000 people from more than 1,300 families were killed inside their homes, underscoring the civilian impact of urban targeting conducted at scale (UN OCHA, 2023).

Taken together, the Gaza case does not demonstrate that AI-DSS inevitably produce unlawful outcomes. Rather, it illustrates how such systems can amplify harm when embedded within permissive rules of engagement, high casualty tolerances, and organisational incentives to accelerate operations. In this

context, AI-DSS function less as tools of refinement than as force multipliers, expanding the universe of targets and compressing decision time in ways that strain the principles of distinction, proportionality, and feasible precautions under international humanitarian law.

Risks and Concerns

The Gaza case study highlights a set of structural risks that arise when AI-enabled decision-support systems are integrated into military targeting under conditions that prioritise speed, scale and operational throughput. While these risks are often discussed separately as technical, ethical or legal concerns, the evidence from Gaza suggests that they are best understood as interconnected failure modes generated by a common institutional dynamic: the compression of decision-making under algorithmically accelerated targeting.

The most fundamental risk concerns the epistemic reliability of AI-DSS outputs. Investigative reporting suggests that systems such as Lavender relied on probabilistic classification models trained on datasets that included civil-defence personnel and other non-combatants whose communication patterns resembled those of militants, reportedly producing an accepted error rate of approximately 10 percent (Abraham, 2024). While such figures cannot be independently verified, the analytical point is broader. AI-DSS necessarily encode the assumptions, categories, and training data selected by their designers and institutions, meaning that any bias or over-inclusiveness at the data or modelling stage is propagated at scale. In targeting contexts, where classification errors carry lethal consequences, even modest error rates become structurally significant.

These epistemic limitations are exacerbated by automation bias. As Cummings (2023) shows, operators under time pressure tend to defer to machine-generated recommendations, particularly when those outputs are framed as statistically grounded or operationally validated. In Gaza, reported practices suggest that human review often functioned as confirmation rather than independent verification, increasing the likelihood that AI-generated suspicions were accepted when they aligned with pre-existing operational assumptions. The risk here is not the absence of human oversight, but its erosion into a procedural formality that no longer meaningfully constrains machine tempo.

A second and analytically primary concern is the acceleration of the targeting cycle itself. Scholars describe this dynamic as characteristic of “hyperwar,” in which data processing and algorithmic decision-support compress the Observe-Orient-Decide-Act loop beyond the limits of sustained human deliberation (Allen and Husain, 2017). The Gaza case illustrates how AI-DSS can multiply potential targets and shorten review timelines, reducing opportunities for legal, ethical or contextual assessment. Crucially, this is not a

technical inevitability but an organisational choice: systems designed to maximise target generation will exert pressure on human operators to adapt to machine speed rather than resist it. Under such conditions, AI-DSS function less as tools of discrimination than as mechanisms for scaling violence.

A third failure mode concerns accountability. As AI-DSS increasingly shape who is flagged, when individuals are locatable and which sites are prioritised, responsibility for lethal outcomes becomes diffused across designers, commanders, analysts and systems. Gounaris and Kosteletos (2020) describe this as a "responsibility gap," in which harm appears to result from technical process rather than human decision. In practice, this risks eroding moral agency and complicating post hoc attribution of responsibility for unlawful strikes, particularly when decision timelines are compressed and auditability is limited.

Finally, the Gaza campaign highlights broader systemic risks associated with the global diffusion of AI-DSS. The systems reportedly used by the IDF resemble emerging capabilities deployed by other militaries, including the United States and Ukraine and are supported by major defence and technology firms. Human Rights Watch (2024) warns that, absent enforceable regulation, such systems may proliferate into conflicts characterised by weaker legal constraints or be exported to regimes that repurpose them for surveillance and repression. The analytical significance of Gaza therefore lies not in its uniqueness, but in its function as an early indicator of how AI-DSS may reshape warfare globally when institutional incentives favour acceleration over restraint.

Policy Recommendations

1. Recalibrate Human Control by Constraining Machine-Driven Tempo

Calls for "meaningful human control" must move beyond formal authorisation requirements and address the pace at which AI-DSS structure targeting decisions. Where systems generate targets at scale, human oversight risks becoming procedural rather than substantive. States should therefore impose tempo constraints on AI-assisted targeting, including minimum review times, mandatory secondary verification for probabilistic classifications, and explicit prohibitions on approval processes that adapt human judgement to machine speed. In dense urban environments, where distinction and proportionality assessments are inherently complex, AI-DSS outputs should be treated as prompts for investigation rather than presumptive indicators of targetability. Meaningful human control, in this context, requires not only human presence but institutional authority to slow, suspend or reject machine-generated targeting streams.

2. Codify Command Responsibility for AI-Mediated Targeting Decisions

To prevent diffusion of responsibility, states should explicitly codify that legal responsibility for AI-assisted strikes rests with identifiable human decision-makers within the chain of command. This requires more than abstract attribution. AI-DSS used in targeting should be subject to mandatory auditability requirements, including preserved records of system outputs, confidence levels, human overrides and approval timelines. Such records should be reviewable in post-strike assessments and, where relevant, legal proceedings. Reinforcing command responsibility in this way directly counters the risk that algorithmic mediation obscures agency or frames lethal outcomes as neutral technical processes.

3. Regulate AI-DSS as Targeting Infrastructure, Not Merely as Weapons

Existing international discussions on autonomous weapons inadequately capture the risks posed by AI-DSS, which influence lethal decisions without independently executing them. International regulatory efforts should therefore treat AI-DSS as targeting infrastructure subject to specific constraints. These should include transparency obligations regarding system purpose and scope, prohibitions on the use of AI-DSS for probabilistic individual targeting without corroborating intelligence, and requirements for ex ante legal review of AI-assisted targeting architectures, not only individual weapons. Without such measures, AI-DSS will continue to escape meaningful regulation while reshaping the conduct of hostilities.

4. Restrict Export and Commercial Enablement of AI-DSS

Given the role of private technology firms in developing and maintaining AI-DSS, export controls must extend beyond traditional weapons to include targeting software, surveillance systems and data-fusion platforms. States should implement licensing regimes that condition export on demonstrated compliance with international humanitarian and human rights law, alongside end-use monitoring and penalties for corporate complicity in misuse. Absent such controls, AI-DSS risk proliferating into conflicts characterised by weaker legal constraints or being repurposed for domestic repression, accelerating the global diffusion of high-casualty, data-driven violence

Conclusion

The Gaza case demonstrates that the principal risk posed by AI-enabled decision-support systems in warfare lies not in autonomy, but in acceleration. When AI-DSS are embedded within permissive rules of engagement and organisational incentives that prioritise speed and volume, they compress decision-making, expand target sets and weaken the substantive role of human judgement. In this context, technical limitations, automation bias and accountability gaps emerge as secondary effects of a deeper institutional choice to adapt human decision-making to machine tempo. While AI-DSS could, in principle, support

verification and legal compliance, the Gaza case suggests that without structural constraints on how such systems are used, they are more likely to scale violence than refine it. As military AI proliferates globally, preventing this outcome requires regulating not only technologies, but the organisational logics into which they are deployed.

Bibliography

Abraham, Y. (2024, April 3). *"Lavender": The AI machine directing Israel's bombing spree in Gaza*. +972 Magazine. https://www.972mag.com

Basuchoudhary, A. (2025). AI and warfare: A rational choice approach. *Eastern Economic Journal, 51*, 74–86.

Birch, M. (2024). Who did that? AI-assisted targeting and the lowering of thresholds in Gaza. *Medicine, Conflict & Survival, 40*(2), 97–100.

Cable News Network. (2023). *Israel's use of unguided munitions in Gaza*. https://www.cnn.com

Cummings, M. (2023). Automation bias in military decision-making. *Journal of Cognitive Engineering*. (Advance online publication if applicable)

Gounaris, N., & Kosteletos, A. (2020). Responsibility gaps in algorithmic warfare. *Journal of Ethics & Technology*.

Gusterson, H. (2016). *Drone warfare: The cultural politics of remote control*. MIT Press.

Gusterson, H. (2024). It's all Lavender in Gaza. *Anthropology Today, 40*(6), 1–2.

Human Rights Watch. (2024, September 10). *Questions and answers: Israeli military's use of digital tools in Gaza*. https://www.hrw.org

International Committee of the Red Cross. (2019). *International humanitarian law and the challenges of contemporary armed conflicts*. ICRC.

Kochavi, A. (2021). *Israel Defense Forces strategic innovation speech*. Israel Defense Forces.

Lowenstein, A. (2023). *Gaza as a lab: Israel's high-tech war* [Podcast episode]. In *TechTonic*. Financial Times. https://www.ft.com

McKernan, B., & Davies, H. (2024, April 3). *"The machine did it coldly": Israel used AI to identify 37,000 Hamas targets*. The Guardian. https://www.theguardian.com

Mhajne, A. (2025). Gaza: Israel's AI human laboratory. *The Cairo Review of Global Affairs*.

Nadibaidze, A., Bode, I., & Zhang, Q. (2024). *AI in military decision-support systems: A review of developments and debates*. AutoNorms, University of Southern Denmark.

United Nations Office for the Coordination of Humanitarian Affairs. (2023). *Hostilities in the Gaza Strip and Israel: Reported impact (Day 45)*. https://www.ochaopt.org

Renic, N. C., & Schwarz, E. (2023, December 19). Inhuman-in-the-loop: AI targeting and the erosion of moral restraint. *Opinio Juris*. https://opiniojuris.org

Opinio Juris. (2024, April 4). *Symposium on military AI and the law of armed conflict: The need for speed, the cost of unregulated AI decision-support systems to civilians*. https://opiniojuris.org

Reichert, R. (2025). Autonomous occupation: Israel's AI-driven drone warfare and the digital architecture of authoritarian power. *Dialogues on Digital Society*.

Sylvia, N. (2024a, February 8). *Israel's targeting AI: How capable is it?* Royal United Services Institute. https://rusi.org

Sylvia, N. (2024b, July 4). *The Israel Defense Forces' use of AI in Gaza: A case of misplaced purpose*. Royal United Services Institute. https://rusi.org

Sylvia, N. (2025). The Israeli military's use of AI in Gaza: Operational efficiency at the cost of humanity. *IEMed Mediterranean Yearbook 2025*.

Thornhill, J. (2023, December 20). *Gaza as a lab: Israel's high-tech war* [Podcast transcript]. Financial Times. https://www.ft.com

Western Sydney University. (2025). The genie is out: Is AI a tool of war or peace? In *Sustainable development without borders*. Western Sydney University.