AI and Modern Warfare: Insights from the 2026 Iran Conflict

The Economy Research Editorial[1,2]

[1]The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

[2]Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

The 2026 Iran war is described as a milestone in the weaponization of Artificial Intelligence on the battlefield. However, in the war the article claims the deployment of AI-enhanced targeting is less of a force multiplier, i.e. a weapon in itself, and more of an agent of change to how military decisions are organized. This war demonstrates that while the increased speed of automated data fusion, automated targeting and a collapsed kill chain might enhance operational tempo, they do so at the cost of increased sub-critical human command. Although the use of inexpensive drones by Iran will undoubtedly prove the continuing relevance of asymmetrical warfare, a more significant war-related conclusion is the emergence of AI-enhanced targeting as the dominant organizational adaptation. Battlefield data can now be viewed as a raw strategic good on which all new decisions of the Department of Defense concerning acquisition, allocation and algorithm improvement will be based. Nevertheless, like highly precise systems of past generations, errors with AI have the potential for being devastating, such as the inclusion of civilians into incorrect targeting databases. In summary, AI-enhanced warfare is the generation of an entirely new problem of accounting: although the authority to use lethal force will remain in human hands, many of the processes surrounding the decision process-information input, rate of processing, selection criteria and operating context-are driven by algorithms and machines. It is thus argued that targeting decisions and rules relating to civilian protection, distinction, proportionality and responsibility of command need to be rethought and designed for the age of algorithmic warfare.

1 Introduction - AI Targeting and the New Logic of War

In the early weeks of 2026, the Iran war marked a decisive escalation in the operational use of Artificial Intelligence in modern warfare.[1] Frequently described as the ”first AI war”, although that label requires qualification.[2] The United States and Israel fielded AI-enabled targeting (a variant of the Palantir Maven platform), capable of fusing large volumes of intelligence data and accelerating the production of targeting recommendations.[3] While the ability of AI to enable approximately one thousand strikes in a single day, an order of magnitude higher than traditional weapon platforms could support and reduce the time it takes from identifying a target to engaging it from hours to seconds all while reducing human judgment in targeting and engagement. Yet this acceleration created serious concerns about error, oversight and civilian casualties.[4] For instance, on February 28, 2026 a strike at Shajareh Tayyebeh, an all-girls school, over 170 people were killed (almost all children). Preliminary reporting suggests that the strike may have involved outdated targeting data in an outdated map entry that the AI incorporated into the target list; a slower and more layered human review process might have increased the chance of detecting the error.[5]

The central argument is that for the Iran war to be properly understood, the integration of AI should be understood not as an evolution of weapons platforms but as a fundamental transformation of warfare.[6] This article argues against the simplistic notion that the important changes are cheap drones versus expensive interceptors and, instead, looks at how AI-driven targeting creates incentives and norms. It will show how military incentives, coupled with the preference of AI for automation will cause meaningful human control to be lost and for an increase in automation bias that creates risks. This could be understood as increasing the mismatch of risks between targets and weapons. This is because, with advances in AI-weapon integration, institutions may come to develop two tiers of warfare: for poorer nations or technologically constrained states, asymmetric warfare built around low-cost drones and missiles, versus for richer nations, data-driven AI defenses based on expensive projectiles. In essence, the data from this conflict becomes a valuable commodity; not only can the US and allies use this data to develop more capable AI-driven interceptors, but adversaries will have a model for imposing costs on a stronger military through cheaper systems.[7]

In the following paragraphs, the article uses quantitative evidence on casualties and cost trade-offs to clarify the policy stakes. The often cited case of using a $2,100,000 missile against a $2,000 drone is perhaps illustrative of why the promise of near-perfect AI-enabled interceptors could push nations to adopt it wholesale.[8] Indeed, a 0.5% error rate against 13,000 targeted drone strikes will still result in 65 civilian casualties. The counter-argument is made: AI can democratize access to military capabilities and improve operational effectiveness. This article shows, based on case studies and quantitative information on the efficiency of AI applications in Gaza and Ukraine (reportedly up to 10% error rate on one Israeli system), that a proliferation of AI capabilities will also result in the erosion of restraint in target selection.[9] Taken together, these considerations present a significant call to re-examine rules of engagement and the laws of war. If machines can select targets on a large scale, accountability and legal standards must follow. Therefore, recommendations have been made that military education must become AI-literate and ethics-based, targeting lists must incorporate review by humans, and states must begin to construct international norms and treaties concerning lethal AI weapons, audit AI-enabled weapons systems, and offer civilian protection through non-lethal defense.

Until such treaties exist, it is crucial to proceed with internal controls and innovation. Among possible internal safeguards is the concept of red-teaming which could be used to search for biases in an AI's target selection.[10] The article will argue for an independent review of especially high-value AI-identified targets. The concept that has emerged among scholars and a few leaders in the military domain is the inclusion of time-lags into the targeting decision to ensure that judgment can overcome technology. If technology rapidly increases the speed of warfare, it must follow that military institutions and legal policies can accommodate it. This war, in short, has provided a preview of the war to come; the question is whether states can absorb its painful lessons before it is too late.

2 AI Is Rewriting the Rules of Modern Warfare

The war in Iran offers clear evidence that it is not merely an ancillary force multiplier but has become central to warfighting operations. Decisions to engage targets once rested on multiple layers of intelligence analysis and approvals. In 2026, this paradigm was rewritten. The U.S. military deployed an AI-powered data fusion system, the Maven Smart System, which synthesized input from satellite reconnaissance, electronic surveillance and drone feeds.[11] In 24 hours during Operation Epic Fury, it processed a data backlog that would previously have required thousands of analysts. According to reporting in the Washington Post, the Pentagon struck 1,000 targets on the first day of the aerial campaign.[12] Combined with language models, the US military went "from being able to hit under a hundred targets a day to being able to hit a thousand." This dramatic five-fold increase in targeting capacity compressed the kill chain: what would have required 2,000 personnel during the 2003 Iraq war to generate 20,000 strikes, now requires a few dozen operators and an AI.[13]

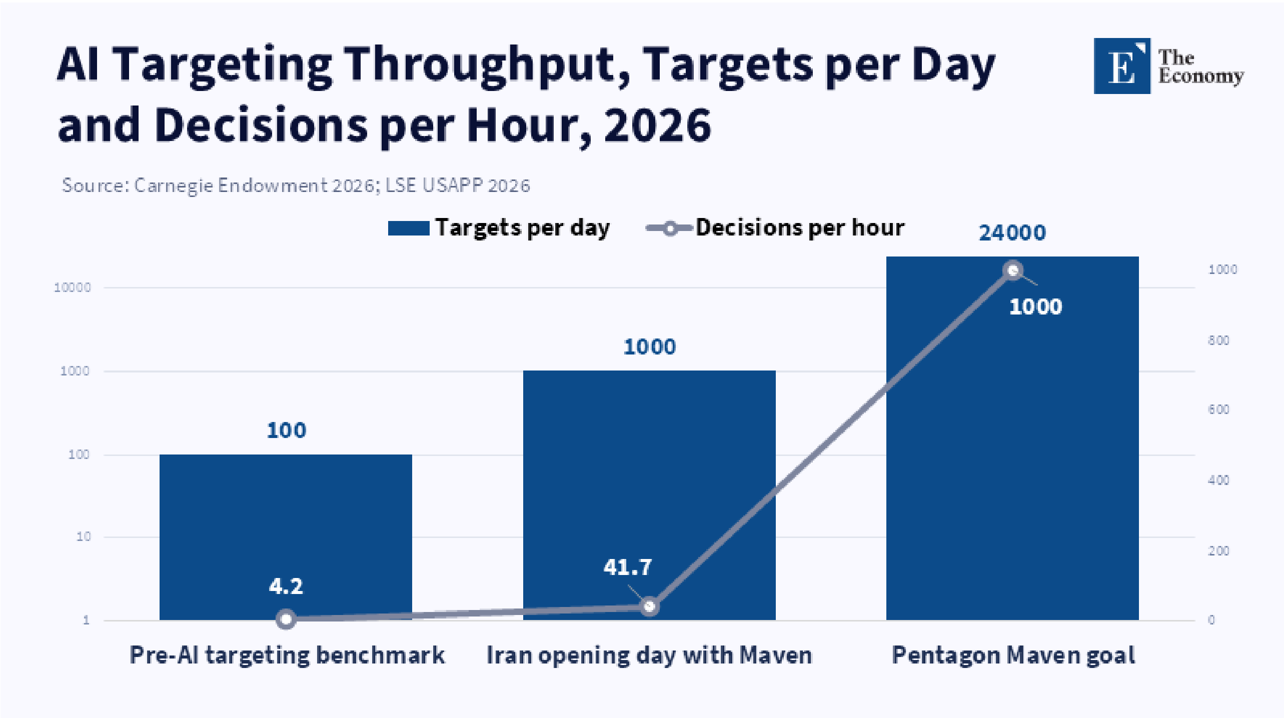

Figure 1: AI-assisted targeting turned the kill chain into a throughput problem, increasing speed while compressing the time available for human review.

This revolution has tangible consequences: the process by which intelligence is translated into actionable strikes has been drastically compressed. Previously, intelligence data might enter a database, then be manually reviewed and approved through a bureaucratic process. Recently, that data package was instantly converted into attack coordinates through a compressed AI-assisted workflow, often while still streaming from original sources (video, radar, radio signals). Naval systems even experimented with autonomously destroying approaching Shahed-type drones. Admiral Brad Cooper (U.S. Central Command) stated that these tools ”allow leaders to make smarter decisions faster than the enemy can react,” shrinking the decision-cycle to seconds.[14] This result is an unprecedented operational tempo: targets that would have required extended review can now be processed and struck within far shorter operational windows.

This dramatic compression of the kill chain has also created a new kind of war signature. The widespread destruction and errors witnessed during the recent conflict, observers suggest, ”raises urgent questions about the accuracy of Maven and the extent to which the Pentagon established proper targeting oversight.” Even supporters argue that when the kill chain is compressed to the point where ”meaningful human oversight becomes largely symbolic,” then ”the balance between speed and accuracy has shifted too far.”[15] Put differently, the core function of AI targeting is to prompt action: if an algorithm recommends a target with 90% confidence, the system is biased toward engagement over further consideration of the 10% uncertainty. This is due in part to the structure of military incentives, which prioritize offensive tempo while treating bureaucratic hurdles (like intelligence validation), which have come to be seen as roadblocks. The devastating cost, analysts note, is that systems focused on efficiency can ”facilitate levels of harm far greater than previously contemplated,” fundamentally undermining principles of restraint in warfare.

Strategic context amplified these dynamics: Iran’s asymmetric attack, tens of thousands of cheap drones required adaptation by Western forces. It brought to the fore private-sector innovation and a relaxation of doctrine: Pentagon software updates during the recent conflict focused on merging AI suggestions with missile targeting guidance and ”automated legal justifications” for strikes.[16] The practical effect is that targeting became a battle of data throughput. American military metrics now emphasize not lives taken or territory gained, but targets processed per hour. For example, by late March 2026, over 13,000 targets in Iran had been struck by U.S. forces including 2,000 command-and-control nodes and 1,450 industrial sites.[17] Yet highly protected civilian assets like Minab’s infrastructure are no longer safe, having been hit by AI-assisted data in the recent conflict for the first time.

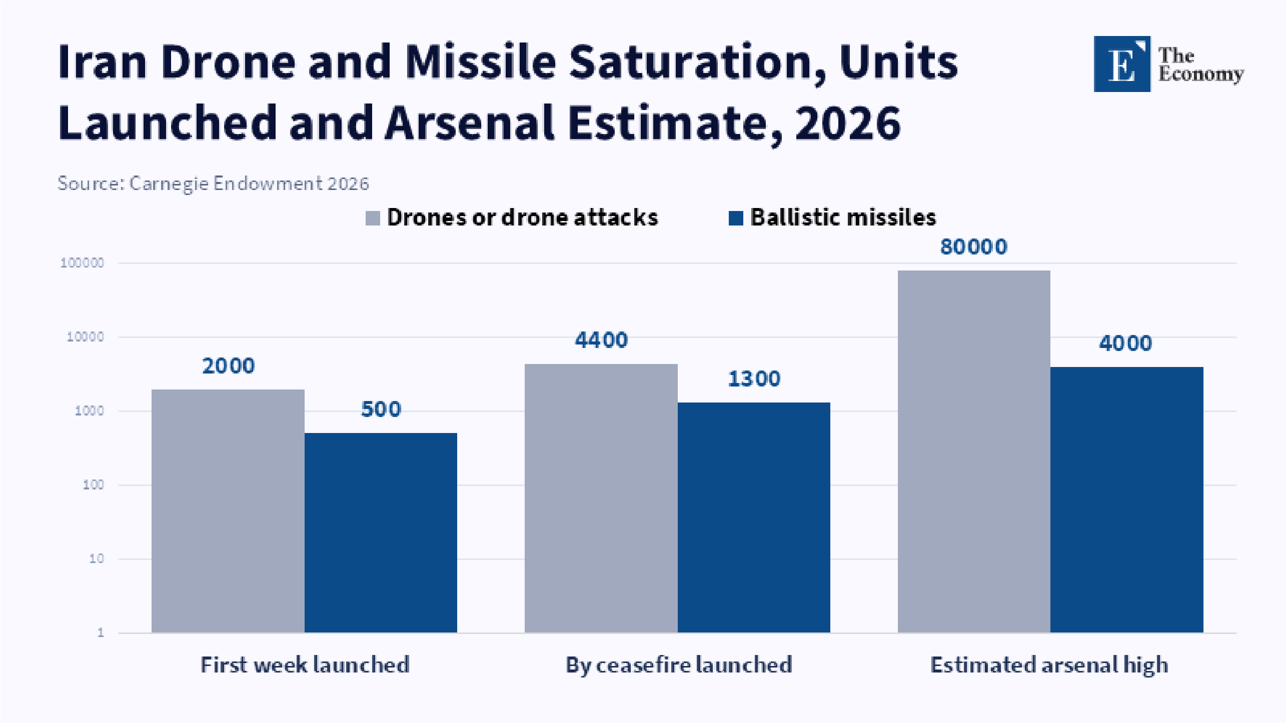

Figure 2: Iran's launch volumes and estimated arsenal illustrate why saturation tactics remain attractive to states facing more expensive defense systems.

A number of comparisons with other recent conflicts can be instructive. While many point to Ukraine and Gaza as previous instances of AI in warfare, their scale and sophistication do not compare to Iran's recent war. Military forces have long used AI to aid analysis and gain situational awareness, analysts concur, but this conflict marked a tipping point where AI-driven targeting became an expectation rather than an exception. Numerous countries including many without advanced technical capabilities, will notice. Iran's asymmetric strategy is seen as a difficult to dismiss gambit and many nations will likely continue to pursue drone swarms and AI weapons as a result.[18] Intelligence agencies believe that a new arms race is underway: one for low-cost swarm tactics, the other for perfecting defenses. The widespread international concern was perhaps best encapsulated by President Zelensky's UN appeal for "global rules now for how AI can be used in weapons."

From this perspective, the ultimate takeaway from the war in Iran is that AI is fundamentally redefining military power and the strategic calculus between humans and machines. The traditional metrics of superiority such as size of armies, number of tanks, are giving way to quantifiable metrics like data assimilation capacity, algorithm sophistication and speed of decision loops. This new definition of military strength will require new doctrines: training will need to prioritize AI literacy and rapid decision-making skills over marksmanship and maneuver formations. Command and control structures, logistics and even industrial base capacity (e.g., warehouses of interceptors versus a few operators using AI) are in flux. In essence, the integration of AI into targeting is collapsing the boundaries between intelligence, targeting and precise strike operations, creating a continuum of conflict that is simultaneously cyber, physical and informational.

3 War Data as Strategic Capital: The Next Defense Investment Cycle

Because the war has been fought at an extraordinary scale and speed, battlefield data has become a valuable strategic asset.[19] The data derived from the thousands of enemy drones that were captured and analyzed, from thermals and strikes that identified target locations and impacts, were logged and examined. According to senior U.S. officials, the weaponization of this wealth of data will only intensify: in the real world, this translates into quicker R&D times and more money allocated. For example, defense spending increases have already been confirmed for 2025-2026. Meanwhile, defense-related venture investment achieved its all-time high in 2025 and has only increased since the advent of the war.[20]

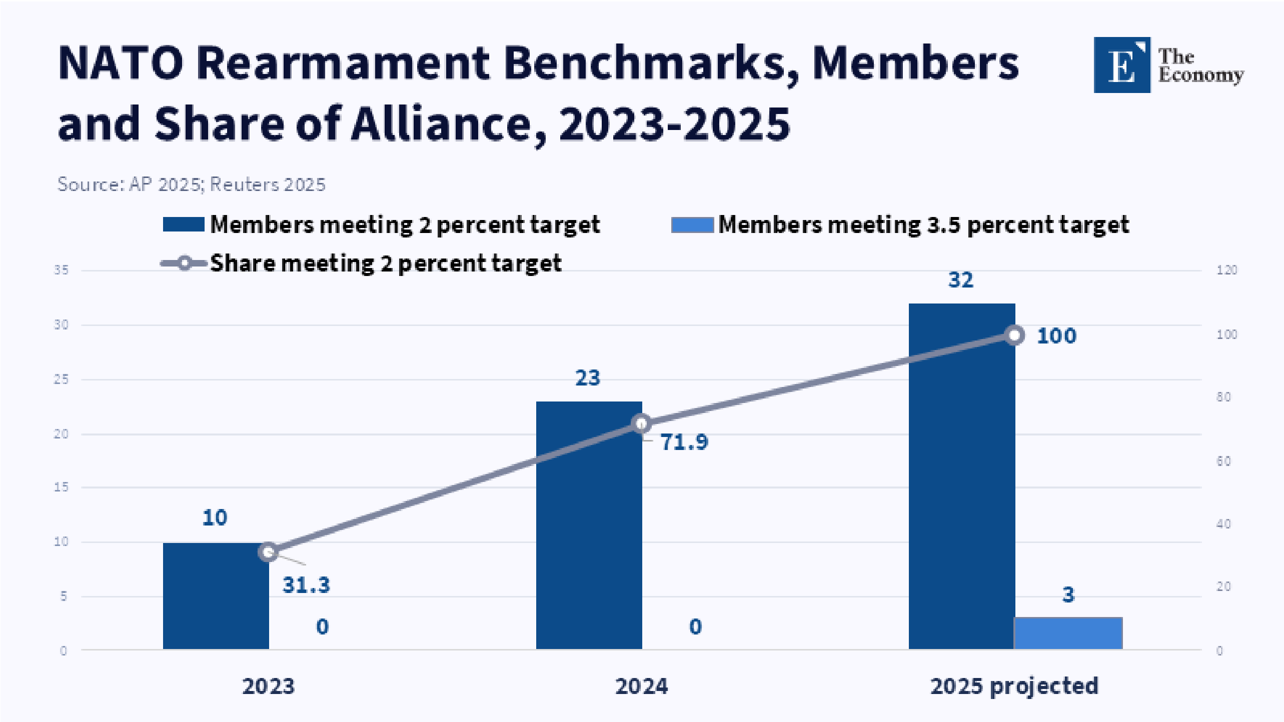

Figure 3: NATO's rising spending benchmarks show how the Iran conflict fits into a wider rearmament cycle that will increase demand for AI-enabled defense systems.

Data has borne out this trend, too. According to one estimate (made by industry data analysis). The U.S. Pentagon budget alone requested tens of millions of dollars for unmanned counter-drone systems in fiscal year 2027 - a product of identifying cheap and viable solutions (such as the LUCAS drone that intercepted Iranian drones); but, as demonstrated, written analyses confirm what Iran discovered in this conflict: unfavorable cost-exchange rates are not tenable. A convenient way to think of this can be shown by Table 1, which compares the costs of standard interceptors against drones of similar capabilities. Numbers from industry analysis display stark disparities- U.S. ships have shot down $2,000 Houthi drones using $2.1 million missiles, a 1,050:1 cost ratio.[21] Western defense systems often cost over $1 million per missile against $20,000 Shaheads.[22] Such imbalanced dynamics will naturally drive future demand for multi-shot, cheap, individual interceptors and AI systems for cheaper single-shot solutions.

2026 Defense System Unit Costs, USD

Weapon/System	Low unit cost USD	High unit cost USD
Neros "Archer" FPV Drone	2000	2000
Shahed/Karrar Drone	20000	50000
LUCAS Drone	35000	35000
AIM-120 AMRAAM	400000	1000000
Standard Missile	2100000	2100000
Patriot PAC-3	3000000	3000000
Block-III UAV / MQ-9	30000000	30000000

Source: Jerusalem Post 2026; Carnegie Endowment 2026; U.S. Department of Defense budget documents

Table 1: The unit-cost gap between drones and advanced interceptors shows why successful defense can still impose a strategic financial burden.[23]

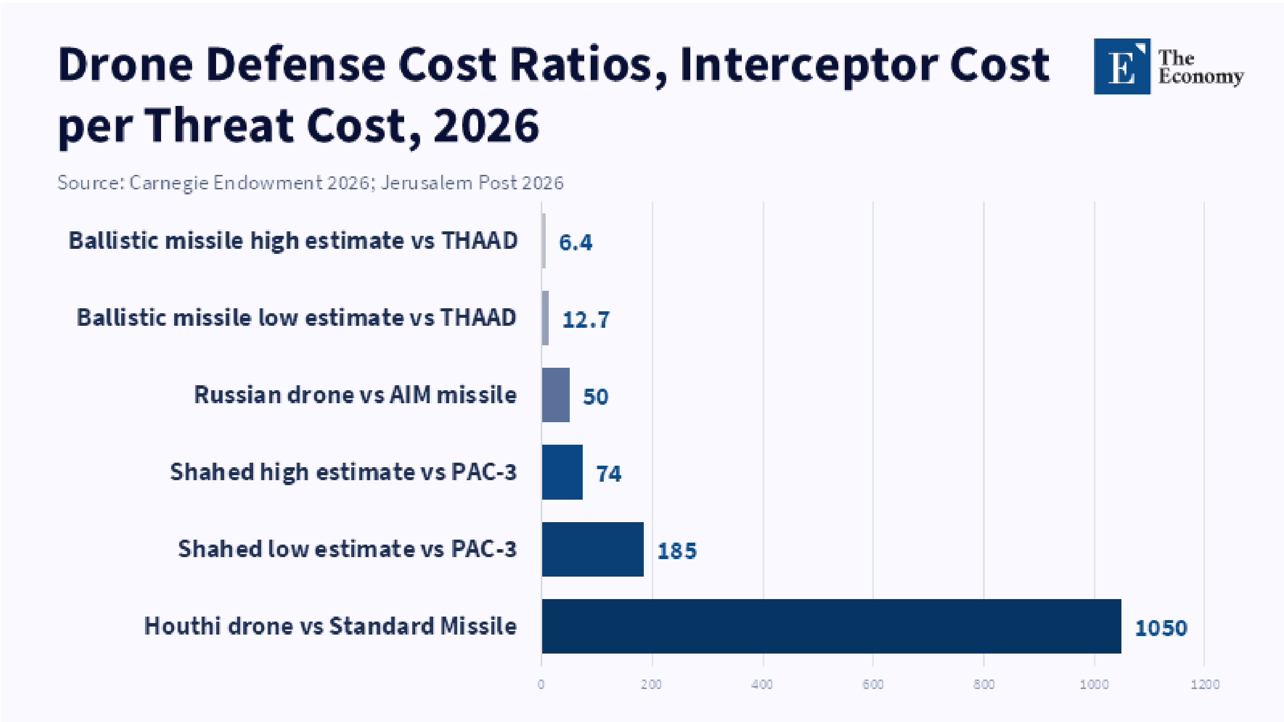

Figure 4: The cost-exchange ratio shows the central asymmetry: defenders may win individual interceptions while losing the economics of repeated defense.

War data became a raw material for the defense industry as well as weapon costs. Companies such as Palantir, Anduril, Shield AI and other defense-technology firms have positioned their systems around lessons from the Iran conflict to stress-test their systems.[24] Others secured sole source contracts to produce and provide further AI modules during conflicts. Congressional testimonies reveal that data from performance within the war in Iran is currently being used by the Department of Defense to refine algorithms and shape priorities in terms of system acquisition.[25] As noted in the Chatham House report, the American military has ”started using its own version of Maven to provide targeting information and NATO also acquired and uses an adapted version.[26]

It seems that a feedback loop is established: the more knowledge obtained from war data, the greater the investment in the development of similar systems. This trend is also evidenced in the conclusion of Carnegie; the world is now facing an accelerating drone and AI arms race and different countries seem keen on emulating low-cost weapons such as the Iranian arsenal or exploiting U.S. AI expertise.[27] This means the future defense budget will increase not only in volume but also in its composition-for instance, pilot and aircrew training might be financed by resources diverted from data scientists or cyber divisions. More broadly, technology companies focused previously on civilian clients will increasingly develop military-grade AI and drones in response to wartime demand.

This trend, however, faces important constraints. A key factor that may moderate the acceleration of uncritical system implementation is the rising costs of errors, which would have diplomatic consequences. If AI systems generate diplomatic crises (as happened when a school strike was investigated), policymakers are likely to introduce constraints; both the militaries of allied countries and the UN received official diplomatic notices during this incident. Such reactions, arising out of civilian harm data, may limit the speed of future investment. Still, the primary tendency is evident: military success is stimulating further investment in war technology. As Nilza Amaral observes in the Chatham House document, ”there remain very real risks”, despite the enthusiasm with which militaries will embrace additional efficiency measures.[28] The pragmatic effect will be a substantial reinforcement of the defense sector’s reliance on AI in offensive and defensive applications, driven by data from Iran. Such development must be coupled with effective post-war assessments, an area that the current organizational structures still have not yet fully developed.

4 From Human Oversight to Automation Bias: The Civilian Risk of AI Targeting Error

The missile could strike civilian housing, resulting in immense loss of life. While the US and Israel have asserted that humans are ultimately in control of lethal decisions, even despite the use of advanced AI, each targeting package did need the approval of the commanding officer.[29] In practice, however, the system was finely tuned to reduce human hesitation and though the Palantir team claimed to use a support tool where there is always a ”human in the loop”, the loop had narrowed severely in terms of length of time for decision-making.[30] In some situations, commanders only had as little as 70 seconds in which to approve each AI-generated strike plan. It is virtually impossible to have an appropriate amount of time for review and now analysts are stating that the system ”inherently biases decisions toward approval.” As a result, the human verification may simply become cosmetic, with officers glancing and approving recommendations to keep pace without hesitation.[31]

The tragic ramifications of such an environment are apparent in the case of a U.S. cruise missile strike on the Shajareh Tayyebeh elementary school on Feb. 28.[32] As investigators have since found, an old military map has failed to recognize that the facility was no longer a military site but had long functioned as a civilian school. The AI system added the target coordinates to the list of strikes and by the time the mistake was discovered, the missile was already en route. Reported death tolls varied but the strike killed a large number of civilians, including many children.[33] The military reported the incident as both "map error" and "human error in wartime"; however, whether the fault was human, data-related or machine-mediated, the outcome was the same. Numerically speaking, this one strike alone accounts for approximately 10% of all children reported as killed in the war so far (254, based on one tally) so despite continuing investigations, this alone is a clear indication that the limitations of AI targeting need to be understood.

Beyond this particular instance, however, patterns of incidents indicate a form of automation bias. Non-government organizations such as Human Rights Watch documented the targeting of dozens of civilian objects (homes, schools, markets, etc.) throughout the war and The New York Times has also uncovered evidence of 22 schools and 17 hospitals that sustained damage and at least 763 schools and 316 health facilities that were hit.[34] This number alone suggests more targets than what could reasonably be considered the result of military necessity and points towards an issue of systematic targeting. There is no numerical data on the actual rate of error for Maven in Iran since the U.S. insists that it is operating as designed; however, other cases offer cautionary comparisons. For instance, an Israeli AI called Lavender in Gaza mistakenly identified civilians as targets around 10% of the time.[35] Even if Maven's error rate were lower, say 1% the problem would remain serious. 1% would yield 130 errors from 13,000 potential strikes which means a missile entering the wrong building or home.

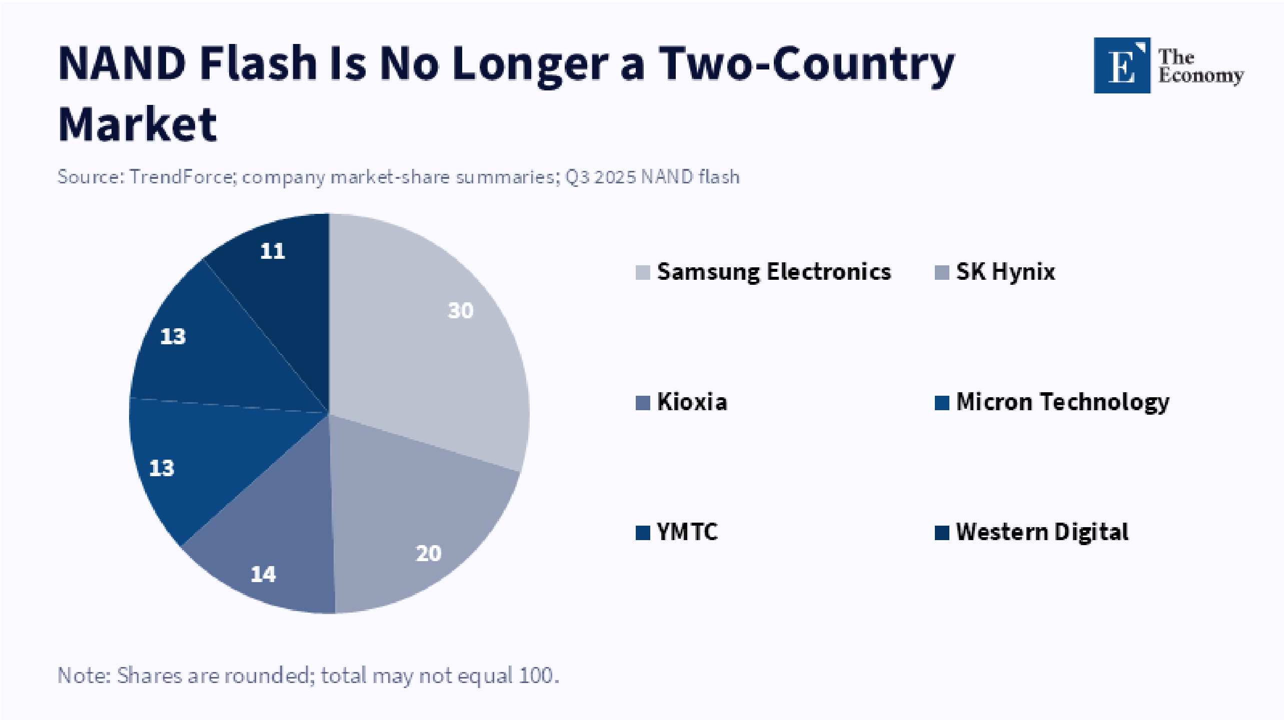

Figure 5: Damage to schools and health facilities shows why AI-assisted targeting must be assessed not only by speed, but by its consequences for civilian protection.

Table 2 details how even public evidence on AI targeting accuracy raises serious questions about civilian harm at scale.

Known AI Targeting Error Claims and Civilian-Harm Indicators

The Economy

Source: Carnegie Endowment 2026; Reuters 2026; Arms Control Association 2026

AI System / Conflict	Estimated Error Rate	Civilian Casualties Linked	Source/Notes
IDF "Lavender" / Gaza 2023	~10%	Thousands of civilians estimated	37,000 suspects reportedly flagged; ~10% error discussion cited in Carnegie analysis
US "Maven" / Iran 2026	Unknown	Large number of civilians; many children	Shajareh Tayyebeh School strike; Reuters reports outdated targeting data may have been involved; total Maven error rate not public

Note: Maven's total error rate in Iran has not been publicly verified. The Lavender figure refers to the reported Gaza system discussed in the Carnegie analysis and should not be treated as evidence of Maven's accuracy.

Table 2: The comparison between Lavender and Maven shows the core uncertainty of AI warfare: documented speed does not yet come with transparent public evidence of targeting accuracy.

This error, in missile mapping, demonstrates a deeper concern: scaling AI does not necessarily mean scaling safety. The more automation, the bigger the mistakes can be. A 0.1% error rate on the targeting algorithm results in 13 massive, erroneous targetings on 13,000 strikes. In a rapid battle setting, erroneous strikes on civilian targets mount extremely quickly. In a pre-automation policy, there might have been twelve different human verification checks in order to strike a busy urban area. Under AI direction, there were none. Many critics fear that with 'shooter-vs-target' logic in machine control (even with a human finally pushing the button), we become further removed from actions and ethical reasoning.

And what about the helpful cases? Advocates of military AI claim that accuracy will far exceed human limitations. The use of big data and pattern recognition should supposedly 'out-spot' camouflaged positions better than a single individual can. The empirical record remains too limited to support strong claims of superior precision. The war in Iran provides one of the first large-scale examples and does not demonstrate accuracy benefits. The Carnegie analysis instead shows that "many destroyed civilian facilities in Iran have explainable justifications suggesting an unsettling concern about Maven's accuracy.[37] One reality is clear: systems will break, but when systems reach their theoretical limits, the law of diminishing returns sets in. The hope that algorithmic perfection will eventually be a software update away has been shattered.

And the future trajectory is therefore troubling. Militaries will not cease to upgrade AI systems. If they can make a system that is 99.9% reliable during testing and the commanders know this, they will trust it more. This was proved in war to be a double-edged sword; overconfidence can easily lead to more tragedy next time. Coordination failures, or enemy deception such as GPS spoofing, will be missed and human lives will be lost in unprecedented numbers far greater than those already visible in Iran today. It becomes clear that the bottleneck must be human control and not the software.[38] As noted in what follows, new legal frames are needed to cope with machine-inflicted wounds as much as we do human malice.

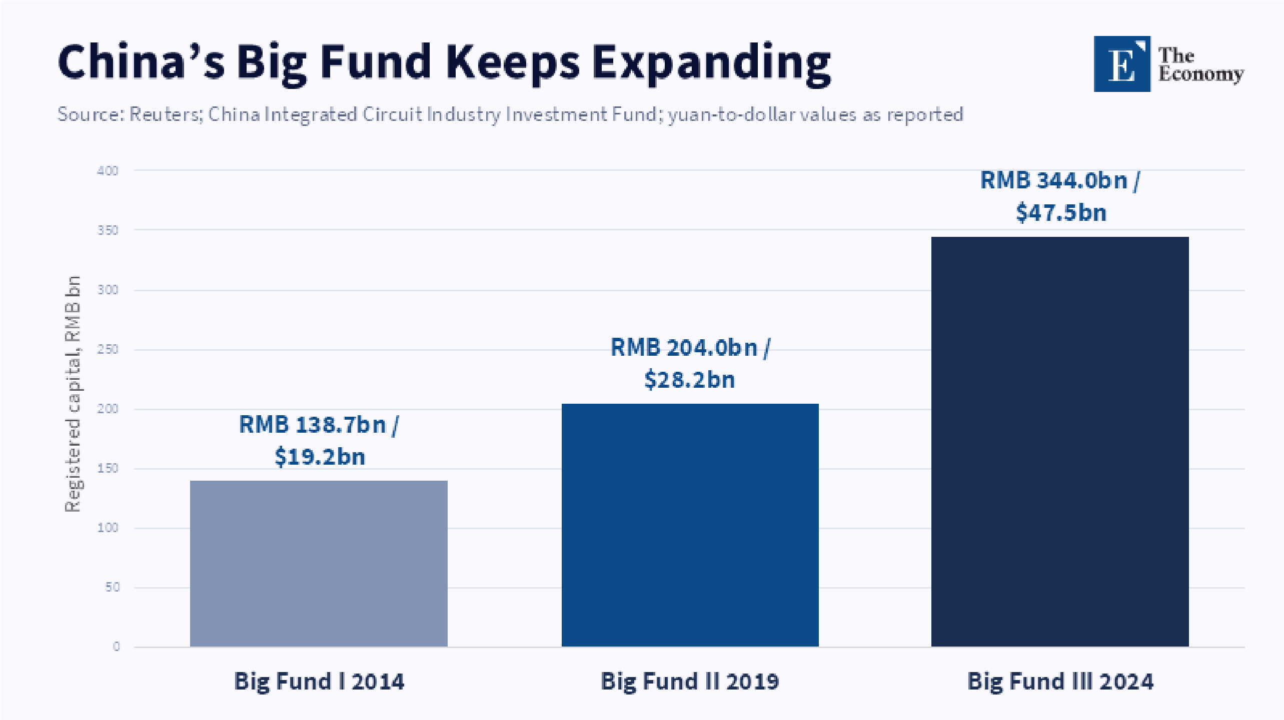

Figure 6: Military AI governance is expanding, but the uneven endorsement of international commitments shows that legal restraint is still lagging behind battlefield adoption.

5 Conclusion - Redefining War Crimes in the Age of AI Warfare

The conflict in Iran suggests that the real danger of AI-powered warfare is not simply machines fighting machines but the gradual, steady transfer of human control and judgment to machines which assert they are, in fact, machine-controlled by humans. Indeed, while automated targeting will accelerate conflict and escalate the scale of target acquisition and battle data processing into future operational benefits, it may also magnify errors, increase machine dependence, and recontextualize civilian injury not as an illegal act or an institutional failing but as a system error.

Thus, the post-Iran debate should ask not only whether AI makes warfighting faster, but whether rules about distinction, proportionality, command responsibility, and civilian protection can survive on battlefields where lists of targets, coordinates, legal rationale, and recommendations on weapons employed are produced and colored by algorithms.[39] After WWII, the international community recognized and developed rules against industrial war crimes.[40] The world of AI-driven war needs similar legal and institutional answers. States should define meaningful human control operationally, include audit trails for AI targeting and enhanced risk reduction measures for civilian harm, and establish who is responsible when algorithmic systems contribute to unlawful damage.[41] Without such safeguards, an accelerating kill chain will be an instrument not of innovation, but the machine-speed erosion of responsibility.

References

[1, 6, 7, 9, 13, 15, 16, 18, 27, 35, 36, 37] Carnegie Endowment for International Peace / Steve Feldstein (2026) *The Unintended Consequences of Iran's Asymmetric Strategy and America's AI War.* Carnegie Endowment for International Peace.

[2, 8, 12, 24, 30] LSE USAPP / Anna Nadibaidze (2026) *US Use of AI in the Iran Conflict Highlights the Continuing Debate Around the Role of People in Warfare.* London School of Economics.

[3, 11, 17, 31] Arms Control Association / Michael T. Klare (2026) *AI Plays Major Role in the War on Iran. Arms Control Today.*

[4, 14, 26, 28, 29] Chatham House / Nilza Amaral (2026) *The Iran War Highlights the Creeping Use of AI in Warfare.* Chatham House.

[5, 25, 32, 33] Reuters / James Pearson and Ryan McNeill (2026) *Bombed Iranian Girls School Had Vivid Website and Yearslong Online Presence.* Reuters.

[6] Vision of Humanity / Institute for Economics & Peace (2026) *How AI Is Rewriting the Rules of Modern Warfare.* Vision of Humanity.

[8, 21, 22, 23] Jerusalem Post (2026) *Winning the Cost-Exchange War in Anti-Drone Defense. The Jerusalem Post.*

[10, 23, 41] U.S. Department of Defense (2023) *DoD Directive 3000.09: Autonomy in Weapon Systems.* Department of Defense.

[19] Jerusalem Post / Merav Davidovits (2026) *Turning Iran's Conflict Data into Defensive AI. The Jerusalem Post.*

[20, 24] Maggie Gray (2026) *How Startups, AI, and Drones Changed the Iran War.* Maggie Gray / Substack.

[23] U.S. Department of Defense (2021) *Fiscal Year 2022 Budget Estimates: Missile Procurement and Aircraft Procurement Cost Data.* Department of Defense.

[34] Human Rights Watch (2026) *Iran Conflict: Civilian Infrastructure and the Laws of War.* Human Rights Watch.

[34] New York Times (2026) *Civilian Infrastructure Damaged in the Iran War. The New York Times.*

[38, 39] International Committee of the Red Cross (2024) *ICRC Position on Autonomous Weapon Systems.* ICRC.

[39] International Criminal Court (2021) *Rome Statute of the International Criminal Court.* International Criminal Court.

[40] EBSCO Research Starters (2026) *War Crimes and Military Justice.* EBSCO.

[41] U.S. Department of State (2023) *Political Declaration on Responsible Military Use of Artificial Intelligence and Autonomy*. U.S. Department of State.