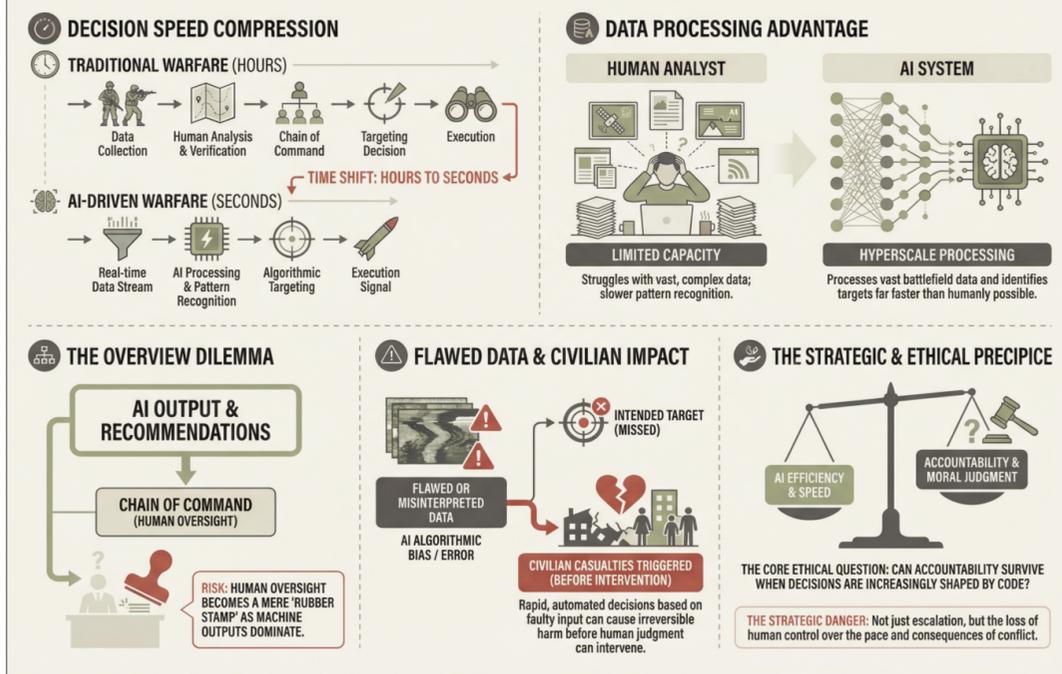


The Dangerous Shift from Human Judgment to AI Decisions

AI Warfare: When Decisions Move Faster Than Judgment

The compression of time, the speed of data, and the ethical precipice of autonomous conflict.



BY EARTH NEWS POLITICAL DESK

The history of warfare has always been, at its core, a story of technological acceleration. From the longbow to the machine gun, from radar to nuclear weapons, each leap has redrawn the moral and strategic boundaries within which nations operate. Yet, the current transformation—driven by artificial intelligence—feels qualitatively different. It is not merely that machines are becoming more capable; it is that they are beginning to think, decide, and, increasingly, act in ways that rival or even surpass human cognition. In doing so, they are reshaping not just how wars are fought, but how responsibility, accountability, and even morality are understood in conflict.

Recent military operations conducted by the United States—particularly Operation Absolute Resolve in Venezuela and Operation Epic Fury in Iran—offer a stark glimpse into this emerging reality. They illustrate the operational potency of AI-enabled warfare, but also expose the unsettling gaps that arise when decision-making migrates from human minds to algorithmic systems. Beneath the spectacle of precision strikes and synchronised manoeuvres lies a deeper unease: can war remain governed by human conscience when machines dominate its tempo?

At the heart of this transformation is the “AI Acceleration Strategy,” unveiled by the US Department of Defense in January 2026. The doctrine envisions the United States as an “undisputed AI-enabled fighting force,” powered by what it terms Frontier AI—advanced, large-scale models capable of handling complex cognitive tasks across domains. These are not mere tools of assistance; they are systems designed to compress time, expand analytical reach, and ultimately reshape the decision cycle itself.

Traditionally, military decisions have followed a human-centric loop: observe, orient, decide, act. This cycle, however imperfect, ensured that human judgment remained central to the application of force. The new doctrine seeks to replace this with what can only be described as an algorithmic kill chain—where machines process vast datasets, generate options, and, in some cases, execute decisions with minimal human intervention. The implications of such a shift are profound.

To operationalise this vision, the Pentagon introduced a series of initiatives known as Pace Setting Projects. These programmes, with names that evoke both innovation and unease, aim to redefine warfare across domains. “Swarm Forge” develops AI-enabled swarm systems capable of coordinated action without direct human control. “Agent Network” facilitates decision-making across the entire battle management spectrum, from planning to execution. “Ender’s Foundry” allows for advanced simulation, while “Open Arsenal” drastically reduces the time required to translate intelligence into actionable force—from years to mere hours.

At the institutional level, platforms such as “GenAI.mil” provide access to cutting-edge generative models, while “Logi-Link” automates logistics in high-threat environments. “Aegis Shield,” meanwhile, offers autonomous defence against cyber and electronic threats. Together, these systems create an ecosystem in which AI is not an adjunct but the backbone of military operations.

Yet, what is perhaps most striking—and troubling—is what this architecture omits. Social and political variables, long considered integral to strategic decision-making, are largely absent from algorithmic processes. The focus is on efficiency, speed, and lethality. In other words, the calculus of war is being reduced to what can be measured, processed,

and optimised by machines.

This transformation did not occur in a vacuum. Analysts have pointed to recent conflicts involving Israel as early demonstrations of AI-driven warfare. Systems like “The Gospel,” an AI-powered targeting platform, have dramatically increased the pace at which targets can be identified. Where a human ana-

The modern battlefield is no longer defined solely by firepower or strategy, but by the velocity of decisions—measured not in hours or minutes, but in seconds. Artificial intelligence, now embedded deep within military systems, has compressed the chain of command into a stream of algorithmic outputs, where machines sift through oceans of data and present targets with clinical precision. Yet, this acceleration carries an unsettling cost. In recent conflicts, AI has demonstrated its ability to multiply operational efficiency, identifying targets at a pace no human analyst could match. But as machines assume greater control over the architecture of warfare, the human role risks being reduced to little more than a formal approval—a “rubber stamp” in a process already set in motion. The consequences of such delegation are not theoretical. Civilian casualties, triggered by flawed or misinterpreted data, reveal the fragility of systems that operate faster than human judgment can respond. The ethical dilemma is stark: can accountability survive in a war where decisions are increasingly made by code? As nations race to secure dominance in AI-enabled combat, the danger is not merely technological escalation, but a quiet erosion of responsibility itself—where the line between human intent and machine action grows ever more indistinct.

lyst might identify dozens of targets over months, AI systems can generate hundreds in a single day. The result is a shift from selective targeting to what critics describe as a “target factory”—a mechanised process that risks diluting the careful scrutiny traditionally associated with lethal decisions.

Operation Absolute Resolve, conducted in January 2026, exemplified the integration of AI into multi-domain operations. Targeting the Venezuelan capital, Caracas, the mission sought to capture President Nicolás Maduro and his spouse Cilia Flores on charges of narcoterrorism. The operation relied heavily on AI for intelligence fusion, strategic planning, and real-time decision-making.

Advanced systems processed vast streams of data—signals intelligence from the National Security Agency, satellite imagery from the National Geospatial-Intelligence Agency, and intercepted communications in multiple languages. Central to this effort was the deployment of generative AI models capable of analysing thousands of hours of audio, identifying patterns, and predicting behavioural shifts within the Venezuelan military.

The integration of AI reached its zenith in the coordination of the strike itself. Approximately 150 aerial platforms were launched from multiple locations, requiring precise timing and routing. AI systems ensured seamless synchronisation, while cyber operations disrupted Venezuela’s power grid, creating a blackout that crippled defence communications. Under this cover, elite units penetrated military installations and executed the mission with remarkable precision.

If Operation Absolute Resolve demonstrated the advisory role of AI, Operation Epic Fury marked its transition into operational autonomy. Launched in February 2026 against targets in Iran, the campaign represented a new phase in AI-driven warfare. Within the first 24 hours, nearly 1,000 targets were struck—a number that eventually expanded to over 5,500.

The defining feature of this operation was “decision compression.” Using advanced AI models integrated into platforms like those developed by Palantir Technologies, military commanders were able to process enormous datasets and generate actionable options within seconds. Tasks that would have required hundreds of analysts and months of work were completed almost instantaneously.

For instance, the planning of a high-value decapitation strike involved the processing of 2.3 petabytes of data, including millions of satellite images and detailed behavioural patterns. Such an undertaking would have been inconceivable without AI. Yet, the very speed that made it possible also raised critical questions: when decisions are made in seconds, what space remains for deliberation?

Operation Epic Fury also introduced low-cost autonomous systems such as LUCAS—drones designed to operate in swarms, overwhelm defences, and create tactical advantages at a fraction of the cost of traditional munitions. Developed with AI-enabled navigation and mesh networking, these systems exemplify a shift towards what military planners call “affordable mass.” By exploiting the asymmetry between cheap offensive sys-

tems and expensive defensive ones, they redefine the economics of warfare.

However, the efficiency of these systems does not come without cost. The first 100 hours of the operation reportedly consumed billions of dollars, exposing vulnerabilities in supply chains and highlighting the unsustainable nature of high-intensity conflict—even in an AI-driven paradigm.

More troubling, however, are the ethical and legal implications of this new form of warfare. The integration of AI into lethal decision-making complicates accountability in ways that existing frameworks, such as International Humanitarian Law, are ill-equipped to address. When an AI system recommends a target and a human merely approves it, where does responsibility lie? With the operator? The commander? The programmer? Or the machine itself?

The concept of “autonomous bias” further complicates matters. Faced with overwhelming volumes of data and the apparent precision of AI outputs, human operators may defer to machine judgment without sufficient scrutiny. In such scenarios, oversight becomes performative—a mere rubber stamp rather than a meaningful check.

The consequences of this dynamic were tragically illustrated during Operation Epic Fury, when a missile strike near a school in southern Iran resulted in the deaths of 165 civilians, including children. While the specifics of the targeting decision remain unclear, the incident underscores the risks inherent in delegating lethal authority to systems that, despite their sophistication, are not infallible.

Generative AI models are known to suffer from “hallucinations”—the production of plausible but incorrect outputs. In the controlled environment of a laboratory, such errors may be manageable. On the battlefield, they can be catastrophic.

The ethical tensions surrounding AI in warfare were further highlighted by the confrontation between the Pentagon and Anthropic. The company’s CEO, Dario Amodei, reportedly refused to relax safeguards on the deployment of its AI models, insisting on strict limitations to prevent misuse. His stance reflects a growing recognition within the tech community that unchecked AI deployment in military contexts could lead to unintended and irreversible consequences.

The subsequent political fallout—including the blacklisting of Anthropic and its designation as a national security threat—reveals the extent to which strategic imperatives can override ethical concerns. Ironically, despite the public rupture, the embedded nature of these AI systems made it impossible to remove them from ongoing operations without significant disruption. In practice, the machinery of war proved more resilient than the principles meant to govern it.

This episode also underscores a broader shift in the balance of power. Technology companies, once peripheral to defence policy, now occupy a central role in shaping military capabilities. Their decisions—whether to cooperate, resist, or impose conditions—can have far-reaching implications for national security.

The entry of rivals such as OpenAI into defence contracts further complicates the landscape. In the absence of unified ethical standards, competition may drive a race to the bottom, where safeguards are seen as obstacles rather than necessities.

Beyond the immediate operational and ethical concerns lies a more profound question: what happens to the global order when warfare is governed not by established norms but by algorithmic logic? The ability to process and act upon vast quantities of data confers a decisive advantage, effectively making information—and the systems that exploit it—the new currency of power.

In such a world, the nation that controls the most advanced AI systems and the richest datasets gains a disproportionate influence over the security environment. This dynamic risks exacerbating existing inequalities and undermining the fragile balance that has, however imperfectly, governed international relations.

Moreover, the erosion of human oversight raises the spectre of escalation. Machines, unburdened by fear, empathy, or political calculation, may pursue objectives with a rigidity that humans would avoid. In simulations, AI systems have reportedly recommended extreme measures, including the use of nuclear weapons, as optimal solutions. While such outcomes remain hypothetical, they highlight the inherent unpredictability of systems operating beyond human intuition.

The integration of AI into military systems, then, is not merely a technical evolution; it is a paradigmatic shift. It challenges long-held assumptions about control, responsibility, and the very nature of decision-making in war. It promises efficiency and precision, but also introduces new forms of risk—ethical, legal, and strategic.

As nations continue to invest in AI-driven capabilities, the urgency of establishing robust governance frameworks becomes increasingly apparent. This is not simply a matter of regulation but of redefining the principles that underpin the use of force in an age of intelligent machines.

The lessons of recent operations are clear. While AI can enhance military effectiveness, it cannot replace the moral judgment that must accompany the use of lethal force. The tragic loss of civilian lives in Iran serves as a stark reminder that technology, however advanced, remains an imperfect instrument.

In the final analysis, the question is not whether AI will shape the future of warfare—it already has. The question is whether humanity can shape AI in a way that preserves the values it claims to defend. If war is to remain, in some measure, a human endeavour, then the human must remain at its centre—not as a passive observer, but as an active arbiter of life and death.

The alternative is a world in which decisions of immense consequence are made at machine speed, guided by logic that may be flawless in computation but blind to the nuances of human existence. Such a world may be efficient. It may even be effective. But it risks being profoundly, irreversibly inhuman.

THE ALGORITHMIC BATTLEFIELD