

Artificial Intelligence Competition as Twenty-First-Century Arms

Racing

The Economy Research Editorial^{1,2}

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

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

Abstract

This paper argues that the competition between the U.S. and China in the field of AI is no longer a race for technological leadership. Rather, it is an arms race of the 21st century, mediated by the civilian infrastructure of chips, cloud services, open-source models, data flows, developer communities, and battlefield feedback loops. Whereas the U.S. While trying to retain a frontier advantage via scale in compute, capital concentration, export restrictions, and militarized integration, China is charting another path via efficiency, the diffusion of open models, industrial penetration, and technocratic statecraft. Recent warfare in the Middle East and Ukraine illustrates why this distinction is crucial. AI is already influencing targeting and intelligence analysis, synthetic media generation, robotic warfare and logistics. I argue that claims that military AI is safe as long as the human "stays in the loop" are misleading; the speed, opaqueness, and machinelike assurance generated can lead to a decline in actual accountability. The crucial policy takeaway is that regulating AI is not a question of market liberalization or of dealing with abstract ethics, but of managing how the capability is transformed into coercive power at specific operational interfaces.

1 Introduction - An Artificial Intelligence arms race in the 21st Century

The race for technological supremacy between the US and China did not materialize in a vacuum. Indeed, it has been set in motion over the last few decades, during the space race and the semiconductor wars. Unlike previous competitions that remained within government laboratories and the corporate world, artificial intelligence is ubiquitous and dual-use. Whereas the space race and semiconductor war did change power relations between states, their core instruments for power production and projection remained relatively static. The true uniqueness of the current AI race, as it currently unfolds, is that intelligence is already saturating our military power, economic productivity, information, and everyday lives.

The reason why it still appears to be a civil race is that there are certain metrics within the civil and military fields that have shown clear development. The global investment in corporate AI grew to \$581.7 billion in 2025, private investment \$344.7 billion, \$285.9 billion of total investment was made in the US alone - a factor of over 23 times the amount spent in China.^[1] In the meantime, the economy-wide adoption of AI reached a record pace among OECD countries, to 20.2 percent of firms in 2025, up from 8.7 percent in 2023,^[2] and the adoption of generative AI had supposedly reach 53% of the general population within 3 years. Not only this, but even the performance gap between Chinese and American models is being erased; The 2026 Stanford AI Index noted that in 2025, Chinese and US Models traded top spots in the performance ranking and that the gap between the leading US and Chinese models had shrunk to 2.7% by March 2026.^[3] Overall military spending, meanwhile, went up to \$2.887 trillion in 2025 - increasing for the 11th year in a row.^[4] The claim is not that every dollar spent on civil AI is immediately weaponized. It is rather that a significant amount of investment in compute power, software platforms, data pipelines, talents, and models overlaps between both domains.

Thus, this paper argues that the apparent race between China and the US to develop the latest AI technologies is more accurately described as a twenty-first-century arms race fought using civil infrastructures and masked in commercial language. The two distinctions of an AI arms race - one on civil infrastructures, and the other of a civil guise - matter both analytically and politically. Thinking of it merely as a technological contest focusing on the rankings of innovation or the benchmarking of models leads to restricted policy options focusing solely on increasing innovation capabilities and model performances. When viewed as a dynamic of arms race, however, our focus must turn toward capability denial, industrial resilience, alliances, data accumulation, precision of targeting, control of information, and other dimensions of military statecraft. AI technology is now perceived as a tool of technological statecraft in the Arab Gulf; Defense, economy, and geopolitics become deeply intertwined.^[5] Such framing is more useful for understanding, as although leading in AI does not automatically translate to global hegemony, it determines, to an ever-increasing degree, a state's capacity to transform economic power into coercive power, and to insist upon other states adopting its architectures, standards, and platforms.

The pressing necessity for this reframing is driven by events occurring in the three-year period 2023-2026. Within these three years, a number of crucial developments occurred. Export control moved from signaling policy to the active denial of critical capabilities. The performance gaps between top US and Chinese models

narrowed dramatically, turning the previous, wide chasm to parity. In addition, open-source Chinese models were distributed rapidly throughout the developer community. Finally, ongoing conflicts are now demonstrating how AI has become a tangible element in military statecraft, from targeting to intelligence analysis and the information domain, moving from the theoretical to the practical. It is no longer a question of whether AI poses military risks; it is instead a matter of how deeply the civil and military applications of a common technology stack are intertwined, and how inadequately policy discourse recognizes the convergence.

2 Different Strategies: US compute denial versus China’s open model proliferation

On the surface, the US and China appear to pursue opposing models. They have drastically different structural endowments; the U.S. still holds the frontier position, having produced 40 cutting-edge models versus 15 in China in 2024,^[6] along with dominance in compute scale, cutting-edge model performance, and access to capital. As reported by the Brookings Institution, U.S. Hyperscalers are expected to spend 650 billion dollars on capital expenditures in 2026;^[7] Chinese firms face capital constraints and restrictions on accessing state-of-the-art chips.^[8] As such, the official U.S. Strategy has been focused on maintaining the edge on the frontier, by developing more scale compute, more infrastructure, and blocking China’s access to critical components, whilst preventing U.S. Hardware and software from assisting Chinese military capabilities. This is the reasoning behind restrictions on chips with advanced nodes, chip-making machinery, and high-bandwidth memory. This policy seems less about pursuing innovation and more like a defensive posture disguised as commercial technology development.

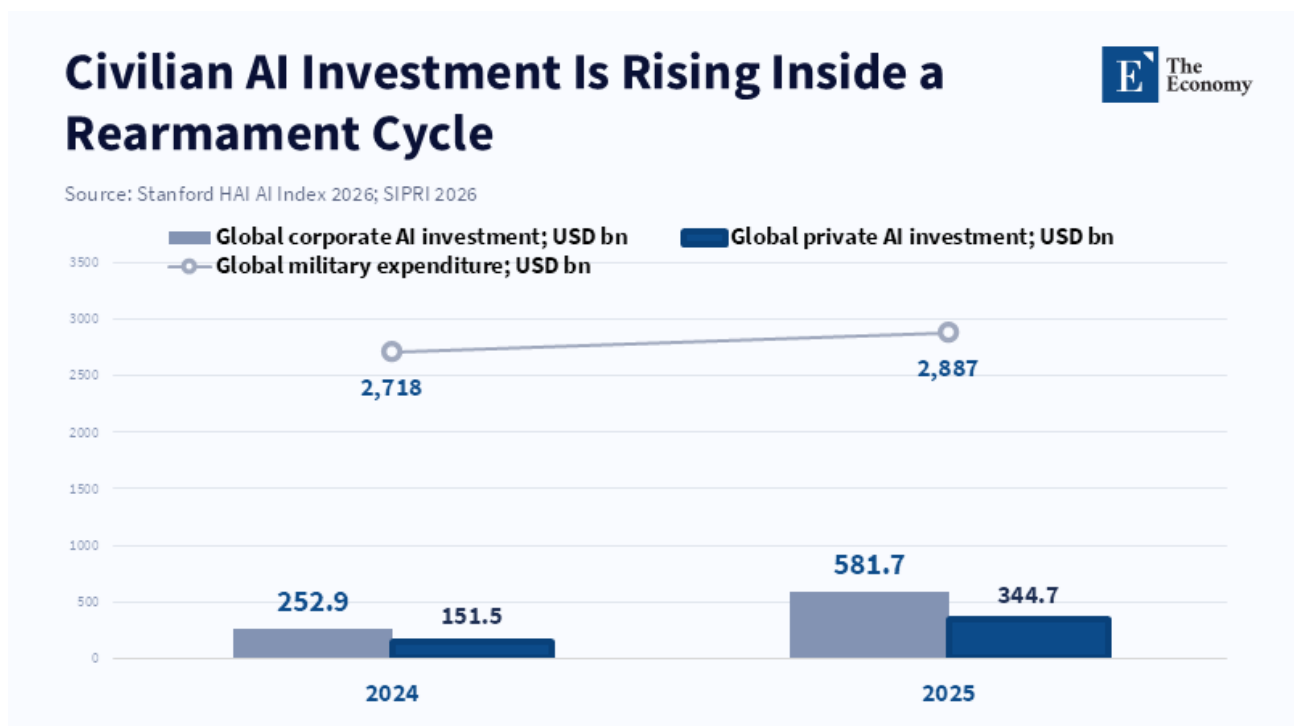


Figure 1: AI appears commercial at the surface, but its investment surge is unfolding during a wider military-spending expansion.

If we strip away the rhetoric, this becomes more apparent. A July 2025 U.S. AI Action Plan talks of the "geostrategic value of open-source, open-weight AI,"^[9] relying on the assumption that the world order will coalesce around whichever ecosystem gains the most adoption. This plan frames access to computing, model proliferation, and the analysis of Chinese frontier models as national security matters, not just business competition. The U.S. Military AI strategy presented in January 2026 offers a more direct articulation of the goals of this policy, defining an AI-driven conflict as a race and laying the groundwork for a fully "AI-first" fighting force.^[10] This plan frames the U.S. Advantage not only in terms of innovation, but also in terms of the data from real-world operations tested during combat. The aim is not to passively observe technology development, but to use existing commercial supremacy as a jumping-off point to gain military supremacy before that edge disappears. The strategy is an indispensable part of an arms race.




Figure 2: The U.S. dominates the financing and company-formation base, while China is much closer in model production.

Although China's tactical approaches differ, its strategic intent is clear. Due to its challenges in winning a pure scale contest (given U.S. Export controls and domestic limitations), China is increasingly focusing on the aspects of efficiency, availability, and distribution. According to Brookings, responses to compute shortage from China are mixture-of-experts models, sparse attention, quantization, and model distillation- basically developing software that tries to extract the maximum compute from limited resources, and is therefore a structurally optimal response to denial. The successes of companies such as DeepSeek and Alibaba are less attributed to achieving top results on all metrics for every cutting-edge U.S. Closed model than on the reasoning abilities, multilingual capability, and distribution cost of their open models. For certain use-cases (especially in regions where proliferation, adaptability, and local control are more critical than top-tier model performance), a low-cost, portable "good enough" model has a significant strategic implication.^[11]




Figure 3: China’s rising share of notable models shows why the race cannot be judged only by U.S. capital dominance.

However, the more significant contribution of China is likely in political economy, rather than model technology. A March 2026 working paper theorized that the open model strategy in China and its manufacturing strengths would become a reinforcing cycle, where open model proliferation drives further use cases and development, and deployment across factories, logistics, robots, research, etc. Provide specialized real-world data, which can be used to further refine the models. In late 2025, China will have surpassed the U.S. in the number of model downloads on HuggingFace, with 7 out of the top 10 downloading models being from China, while the total downloads of derivatives of Alibaba’s Qwen exceeded 100k downloads. If, the paper argues, the source of value creation in AI moves away from cutting-edge training to the deployment of low-cost open models for general-purpose use cases ranging from factories, logistics, research, to robots, then a nation will have a lasting strategic advantage even with a smaller compute scale on the frontier. The notion that an arms race will be won not by training the biggest model, but by the largest distribution of adaptive AI in the economy merits strong consideration.^[12]




Figure 4: Frontier performance is no longer separated by a wide gap, pushing the strategic race toward cost, diffusion, and deployment.

This is being accompanied by a broad geostrategic narrative. A 2026 study from RAND shows that out of the web traffic to LLM websites in 135 countries, 93% was from the U.S. In August 2025, this figure had dropped rapidly, from 3% to 13%, in merely two months, after the release of DeepSeek R1. Over 10% market share is achieved in 30 countries in this metric, and over 20% in 11 countries, particularly in the developing world and countries that trade heavily with China. According to RAND, these gains are attributed to affordable prices, multilingual support, and AI diplomacy. China’s LLM prices are about one-sixth to one-fourth of those from U.S competitors.^[13] One peer-reviewed paper on the Arab Gulf shows a more abstract description of this dynamic; the U.S. utilizes military AI to bond allies to its security structure, while China utilizes technology diplomacy to create a sphere of influence; “soft power” is reimagined as dependence on external models for capabilities.^[14]



Figure 5: China's open-model ecosystem turns adoption, reuse, and developer dependency into geopolitical leverage.

One may object that this is only a form of conventional competition: firms invest, technology advances, users benefit and the costs of deployment are reduced. This has a substantial basis of truth. An OECD report of early trends suggests that AI is capable of increasing productivity substantially in particular tasks and in terms of initial users.^[15] Stanford's 2025 AI Index reports that 78% of organizations employ AI, while the cost to run GPT-3.5 equivalent has dropped over 280-fold between late 2022 and late 2024.^[16] The potential of AI to provide actual benefits must not be discounted. However, it is in the context

Due to its increased affordability, generalization, and usability, AI's strategic implications grow. Calculators have not induced export controls, and search engines did not produce military AI strategies to achieve an "AI-first" doctrine. The fact that open model proliferation, compute access, and deployment architectures have now been translated to national security concerns means that the narrative of productivity gains and the narrative of an arms race are not contradictory: one represents the setting of the other.

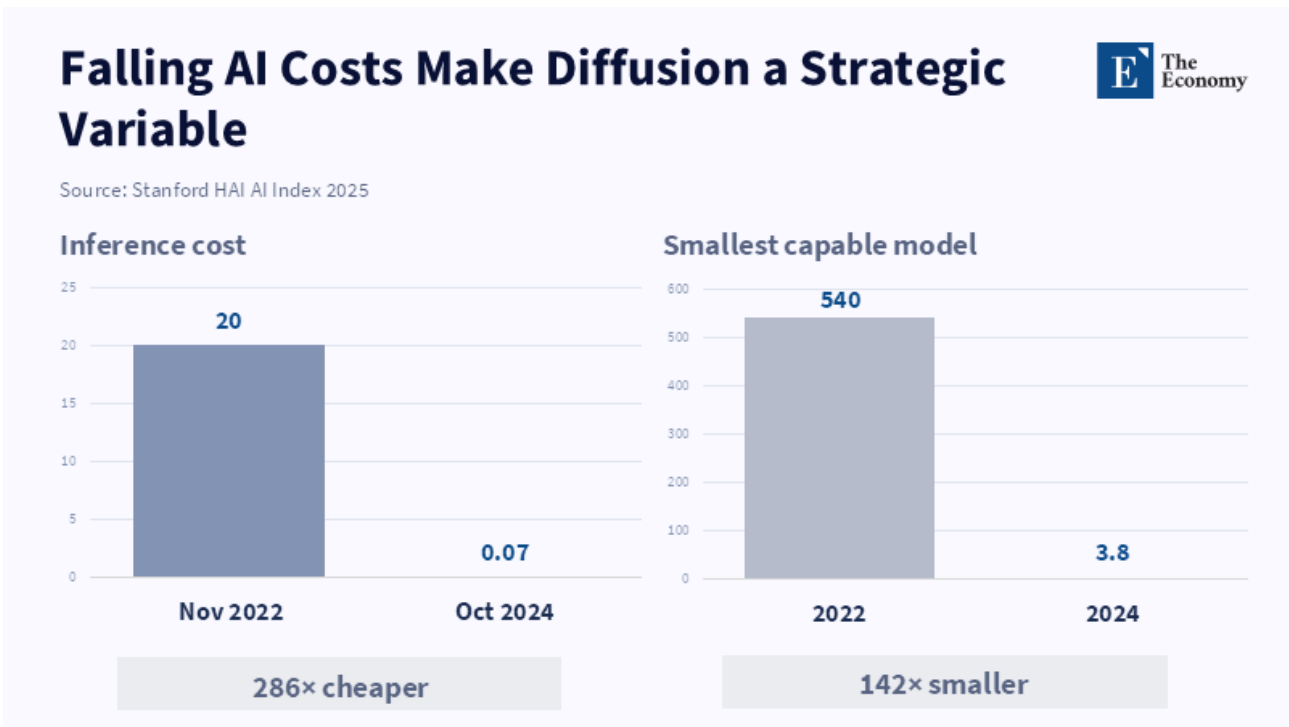


Figure 6: Cheaper inference and smaller capable models make AI easier to embed across industries, infrastructure, and military systems.

3 AI in Warfare: The Invisible Military Logic of the Tech Race

It is precisely in current battlefield operations that the strategic meaning of these competing development strategies becomes most evident. As reported by CBS in March 2026, the U.S. Military was examining over 1,000 potential targets per day, striking most of them, with the time until the next engagement potentially reduced to less than 4 hours. According to former officials, AI is that which can process the sensory inputs from a battlefield inundated with images, documents, feeds and videos, far faster than any team of human analysts could; "It is not the AI actually doing the flying of the planes or launching of the missiles, that will remain a human action. It's everything that comes before-the construction of a targeting package, the assignment of assets, the shortening of battle times from days to hours." This is a crucial distinction. AI's military contribution is not one of autonomously deciding and carrying out attacks but rather, serving as an intermediate layer shaping what commanders "see" and "do not see," how fast they "see it" and when that image has presented them with an "actionable option" at the point of combat authorization.^[17]

That same transformation in operational logic appears in formal military planning strategies from both sides of the Pacific. A Chinese study published in March 2026 on PLA writings describes the incorporation of AI into "intelligentization" and details four objectives across Chinese kill-chains: cross-domain information integration, integrated functionality, multi-agent collaboration, and fast data processing.^[18] The January 2026 U.S. Military strategy document seems to follow this same intuition through different rhetorical language, emphasizing leveraging U.S. Strengths in AI computing, model leadership, entrepreneurial innovation, and operational data to integrate AI throughout the total force.^[19] These converging developments are strategically significant: they indicate that the military value of AI is not that it will soon facilitate the use of an all-

encompassing robot war, but rather that it can fundamentally re-structure intelligence gathering and analysis, command and control, logistics and planning to enable an increasingly software-defined management of warfare. When both adversaries operate with this as the central assumption of military strategy, then claims that military AI development is only a collateral effect of the civilian race appear highly unconvincing.

The strongest argument for military AI suggests that this transformation will not be carried out irresponsibly. According to this perspective, the use of AI in target selection and identification is not synonymous with handing over to machines the authority to make kill decisions. Decision-support systems, when closely supervised, can accelerate identification of patterns, streamline the discovery of likely targets, and ultimately aid humans in upholding the laws of armed conflict. Legal interpretations like that from the Lieber Institute describe systems like Lavender and Gospel as decision-support systems, rather than decision-makers, and argue that they will actually serve to enhance target identification, with a human operator ultimately responsible.^[20] On paper, this argument is highly persuasive: it is true that modern warfare generates data that humans cannot possibly process in time, and there is a genuine argument to be made that to defend cities from mass missile attack requires machine-level targeting priorities in order to protect civilians. An in-depth examination of this complex question must give full credence to these claims.

The problem, however, is that even when a "human in the loop" remains formally present, this phrase can easily become substantively hollow within an organization. Carnegie explains that "instead of piercing the fog of war with clarity, AI is in fact creating a new fog of war, built on the speed, opacity, uncertainty, and lack of accountability that characterize machine intelligence."^[21] A Chatham House analysis of the 2026 Iran war notes that U.S. Commanders confirmed the use of highly advanced AI systems to process unprecedented volumes of data and dramatically shorten the decision cycle for lethal targeting to minutes rather than hours or days.^[22] Even with the human component, there remains the crucial question of whether time spent analyzing AI-provided options in a few seconds rather than a few hours is sufficient for genuinely independent thought and critical questioning when presented with the output of a pre-prioritized list of likely targets. The point, therefore, is not simply whether a person agrees to the strike, but whether the system's dynamics accommodate significant dissent, scrutiny and contextual awareness within the temporal constraints of rapidly unfolding warfare.^[23] Nominal human control in a high-speed combat scenario can simply devolve into a superficial post-hoc validation of machine-generated decision-making.

This, above all, is why the legal and ethical debate surrounding AI's role in war has begun to evolve from solely assessing autonomy to evaluating the design of AI-enhanced decision support. SIPRI's reports from 2025 frame the comparison clearly: comparing autonomous weapons and AI-driven decision support for military target selection, they emphasize that in both cases, human influence on force decisions is being changed, leading to policy gaps within the current legal and ethical frameworks governing warfare. In the same vein, a SIPRI report focuses on bias within military AI and explains how bias, introduced at any point from data input to modeling and deployment context, may well lead to violations of the principles of discrimination, proportionality and precautionary measures.^[24] It is not necessarily runaway artificial intelligence, but the everyday failure of human institutions to adequately account for system errors-e.g., inadequate training, implicit biases encoded in algorithms, over-reliance on outputs, poor monitoring of recommendation generation-that pose the threat.

A September 2024 report on digital tools in Gaza revealed how many of those used had faulty data and crude approximations that could cause unlawful civilian casualties.^[25] The structural vulnerability is machine-dependence under conditions of speed, opacity and weakened accountability.

Recent conflicts are not only evidencing AI’s impact on the physical, kinetic aspects of war; they demonstrate how AI has become a tool in a much broader battle over the information environment. The aforementioned Brookings analysis of the 2026 Iran conflict observed a massive proliferation of fabricated content on social media, ranging from falsified images and battle damage to synthetic video clips created by commercial-grade generative tools.^[26] By analyzing Community Notes data from X,^[27] the report determined that since the conflict’s commencement, over 5,000 such notes have appeared in the digital sphere and over two months preceding the report, posts referencing AI-generated content accounted for approximately 7-8 percent of all disputed content, while AI-generated videos received millions of views during the initial two weeks of fighting. This data is not marginal to strategic considerations. Shaping the information environment is critical for crisis signaling, public perception, alliance coherence, and state credibility when claiming success or failure on the battlefield. Inasmuch as AI can now degrade the information environment with wartime-level speed, it inherently augments the coercive power of the state, irrespective of whether a machine is directly involved in firing shots.

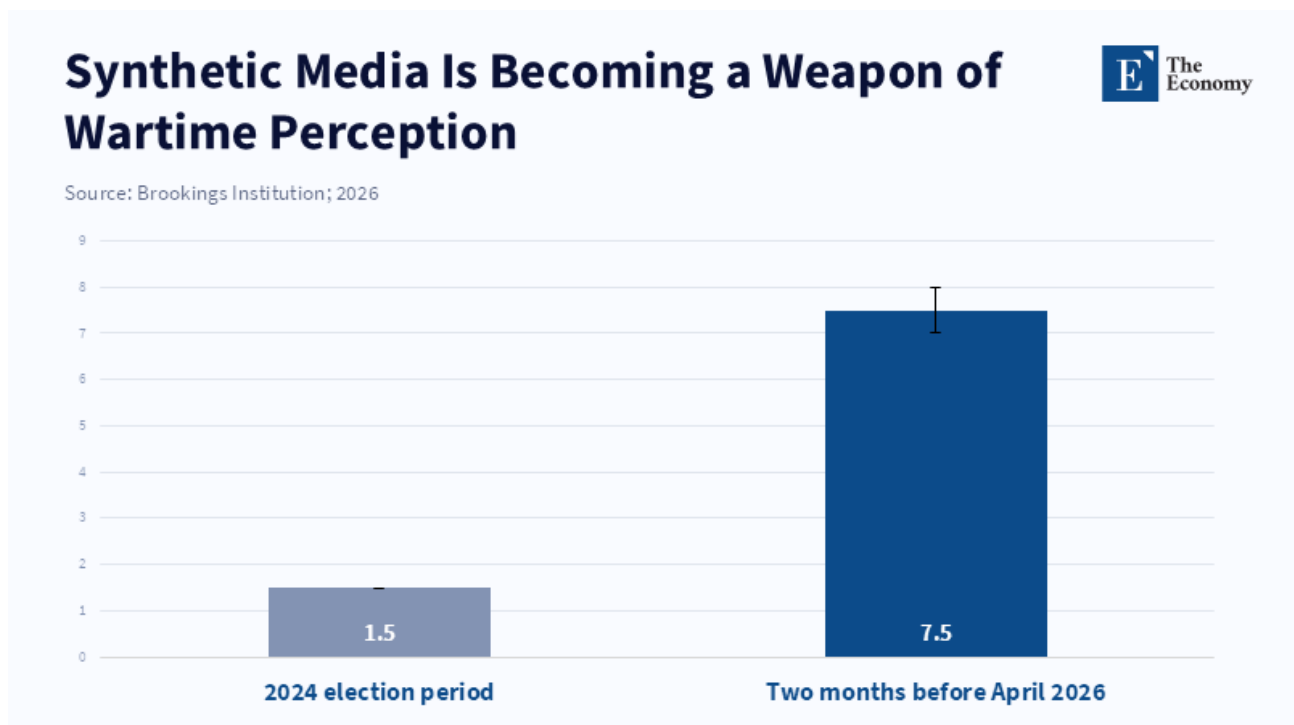


Figure 7: Synthetic media does not need to dominate wartime information to weaken verification, trust, and crisis interpretation.

The March 2026 call for a moratorium on military AI by Nature, while understandable from a moral standpoint, is not a strategically viable policy option given the current context.^[28] Militaries globally have already begun to integrate AI into targeting support, surveillance data processing, unmanned aerial vehicle navigation, and operational management. A broad moratorium lacks a realistic compliance framework, given the incentives confronting competitors actively involved in combat, and has no way to effectively enforce a halt in weaponizing artificial intelligence in the face of potential strategic advantage by not doing so. This

doesn't make policy solutions impossible, but rather highlights the need for a more nuanced approach than an outright ethical ban. Policies must be specifically operational: they need mechanisms for auditing targeting recommendations, traceable data chains, robust testing of AI systems in real-world scenarios, separate legal frameworks for AI-enhanced decision support systems, and institutional designs that retain genuine scope for human dissent when the pace of warfare begins to accelerate beyond human speed. NATO's recently updated AI strategy prioritizes testing, evaluation, validation, verification, certification, and data quality, which aligns with this more practical way forward.^[29]

4 Real World Examples: The Middle East and Ukraine as AI battlefields

The military-strategic logic of AI competition is most evident in the Arab Gulf. An unpublished 2025 study suggested that military use of AI must be viewed as an aspect of technological statecraft, "as the battlefield becomes, simultaneously, both the site of conflict and the primary arena of strategic competition for mastery of cutting-edge technological means in a multipolar world."^[30] In this part of the world, US security support for its Arab partners remains central while China has become the leading trading partner of several Gulf states. AI will likely serve to mediate between these two relationships and allow the states concerned to reconfigure them. The United States will rely on "the U.S. Military AI advantage and defense partnerships to lock allies to the U.S. Orbit in security." China, meanwhile, may leverage its "inexpensive model ecosystem, connections to infrastructure and openness to technology to expand its reach in the region without directly duplicating US security relationships." As such, the theoretical and practical implications of the Gulf case for the growth of AI competition are profound. It is not simply about more accurate, more powerful AI models but about the system of governance to which other states may increasingly be beholden.

The case of Gaza may help to better illustrate this difficult transformation. An investigation that appeared in September 2024 suggested that the Israeli military relied on faulty data for attack decisions in Gaza and thus ran the risk of violating international humanitarian law.^[31] This is relevant not because the AI system utilized in this instance was autonomous but because it may be an indicator of how digitized systems routinize a shift in the locus of responsibility. The presence of a network of sensors, machine learning tools, and AI-assisted targeting systems does not necessarily absolve human decision-makers, but it does distribute the locus of decision-making more widely. A human may ultimately authorize the use of a weapon but it is a human whose decision is guided to an important degree by information processed by a machine and sometimes by machines whose limitations are not fully apparent. The argument that the use of these systems could increase compliance with humanitarian law, so long as they are carefully monitored and employed, is valid, but in rapid, mass, or high-tempo conflict, speed of action is given priority over consideration of existing knowledge.

In this context, this shift to a new domain was reflected in the kinetic and informational aspects of the March 2026 conflict between Iran and Israel. The report from Brookings pointed to an "exponential increase in false and synthetic wartime content"^[32] while Chatham House expressed concerns over AI-assisted targeting after

the April 2026 attack on a girls' school in Minab.^[33] In this situation, U.S. Military commanders themselves admitted to the use of highly sophisticated AI tools to help analyze data and to accelerate decision-making. Here, the question is not simply the prospect of an AI-generated error but also "the convergence of machine-assisted targeting with machine-assisted disinformation in a single conflict." In this environment, the AI system will reduce the time from sensor to target and from target to fire while extending the time from information to truth. Both allow for a higher tempo of operations and reduce accountability.

This "bloodless war" concept associated with AI is, therefore, misleading; it is indicative of a failure to recognize this transformation. Though war may become remoter for the state employing the AI because there are fewer casualties among their soldiers, fewer sorties are necessary thanks to drone reconnaissance, weapons have longer ranges, and domestic political costs of war have been minimized, it does not follow that the killing involves less bloodshed. For those on the receiving end of fire, bloodshed may be continuous and scalable. More generally, the ability to remove oneself from direct involvement in combat diminishes domestic political sensitivity to war, while accelerated, high-velocity information warfare may simultaneously diminish the citizenry's ability to assess what has truly been occurring on the ground. Far from removing bloodshed from the war, AI may be reshaping where and by whom killing takes place. That is ultimately why, more than ever, mastery over AI is becoming a prerequisite for strategic primacy.

The case of Ukraine pushes this trend even further into focus. In February 2026, Reuters reported that much of the nearly 1,200 km front was dominated by unmanned first-person-view drones, making the war a "de facto air war of mutual denial and increasingly resulting in lethal impacts."^[34] Drone-induced casualties increased from below ten percent of the total in 2022 to as high as eighty percent in 2025. These events led to unprecedented difficulty in front of assaults, extreme danger in everyday movement, and evacuation times lasting up to three days in many areas. Although unmanned ground vehicles were widely used for logistics and evacuations (Ukraine alone logged more than seven thousand such missions in January 2026),^[35] the overall picture reflects not just a technological shift but a fundamental reorganization of war in favor of cheap sensing, cheap precision, and constant human-machine interaction facilitated by software.




Figure 8: Ukraine shows that the AI arms race is not only about frontier models, but about battlefield data, robotic logistics, and rapid software iteration.

Ukraine also exemplifies how war has become a data source. By March 2026, it was sharing battlefield data with allies and the private companies helping to build its drone AI, yielding millions of labeled images from tens of thousands of flights.^[36] Earlier in the year, Reuters reported, Ukraine was already using AI to support its drone pilots and to analyze satellite and aerial imagery at speeds that would have taken human analysis days. Ukrainian officials project the output of approximately four-and-a-half million drones in 2025^[37] and are testing them as hundreds or even thousands of miles distant interceptors. Ukraine is demonstrating not only its innovation but also how the battlefield itself is becoming an information generator that will shape AI development, acquisition plans, and the nature of technology transfer between allies. The war is transforming into a live training facility for AI.

Beyond Ukraine, this iterative feedback loop informs experimentation elsewhere and around the world. NATO and Romania were observed in April 2026 testing AI-assisted drone interceptors near the Black Sea, and their operational deployment was imminent.^[39] The Reuters report indicated that the United States would also begin supplying drone interceptors to the Middle East to protect its allies from Iranian attacks. This is clearly not simply an instance of consumption of U.S. Training and hardware by an ally, but also of export of battle-tested data-centric approaches, software-enabled command-and-control capabilities, and operational doctrine by Ukraine. Classic arms races rely on observation and imitation; AI arms races are based on data analysis, flexible acquisition and rapid iteration, and continuous interaction between human and machine users on the battlefield. While advanced models are still critical, the rate of iteration will increasingly become a defining aspect of relative power.

These examples highlight the necessity and inadequacy of export controls. Denying compute power will buy adversaries time at the frontier, but significant strategic adaptations will come more from the integration

of smaller models, distributed systems, robotic hardware, operational data, and battlefield usage than from silicon fabrication or large model speed. Nations must therefore develop alternative measures to silicon volume and benchmark performance to track national advantage, and index factors such as deployment density, the secondary model ecosystem, production integration, and data availability. Allies should develop common testing and certification bodies for military AI, and maintain separate multilateral talks concerning AI-powered decision support (as distinct from less significant issues such as autonomy). Procurement will need to incorporate a traceable pipeline for targeting suggestions, auditable records for review of decisions, continuous bias tests, and a verifiable flow from data intake to targeting recommendations. Wartime use of synthetic media will need to be considered a national security matter (not just a platform moderation issue).

5 Conclusion - Governing Artificial Intelligence as a Catalyst for Strategic Power

The underlying falsehood of the common notion of a U.S.-China AI race is its description as a civilian technology competition that might ultimately translate into a military advantage. Instead, this is already an arms race based on chips, models, data, platforms, logistics and the actual use of AI on the battlefield. The U.S. Approach to maintaining an advantage at the frontier with compute, capital, export controls and alliances will likely be met by China's use of efficiency, wide access to open models, integration of the technology into existing industries and technological statecraft. These trends-evident in the conflicts in the Middle East and Ukraine-can not be considered theoretical: they represent faster targeting, a redesigned information war and the militarization of battlefields in the context of live data.

The government must now stop viewing AI solely as a market technology with potential military applications or as an abstract notion in ethical discussions without concrete consequences. While export controls still matter, they should be supplemented by a combined approach that includes testing and certification standards among allies, auditable military procurement systems that include AI under the right to test, verifiable data systems and non-discriminatory rules for AI-assisted military decision-support. The actual questions to ask now are not which side has a better model, but who manages the conversion of capability into power within current frameworks, and who will govern its use as a coercive instrument? AI has already become integrated into systems of coercion and can no longer be denied as a political strategy, an institutional necessity and a democratic obligation.

References

- [1, 3] Stanford Institute for Human-Centered Artificial Intelligence (2026) *Artificial Intelligence Index Report 2026*. Stanford University.
- [2] OECD (2026) 'AI use by individuals surges across the OECD as adoption by firms continues to expand', *OECD*, 28 January.
- [4] Stockholm International Peace Research Institute (2026) 'Global military spending rise continues as European

and Asian expenditures surge’, *SIPRI*, 27 April.

- [5, 14, 30] SAGE Journals (2025) ‘Military AI, technological statecraft and strategic competition in the Arab Gulf’, *European Journal of International Relations*. doi:10.1177/03043754251377866.
- [6, 16] Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarasci, L., Lotufo, J. B., Rome, A., Shi, A. and Oak, S. (2025) *Artificial Intelligence Index Report 2025*. Stanford Institute for Human-Centered Artificial Intelligence.
- [7] Williams, S. (2026) ‘Hyperscalers Are Spending Nearly \$700 Billion in 2026 on AI Infrastructure — but This Pales in Comparison to the Estimated \$1 Trillion Spent by S&P 500 Companies on Another “Growth” Initiative’, *The Motley Fool*, 16 March.
- [8, 11] Chan, K. (2026) ‘China is running multiple AI races’, *Brookings Institution*.
- [9] White House (2025) *America’s AI Action Plan*. Washington, DC: Executive Office of the President.
- [10, 19] U.S. Department of Defense (2026) *Department of Defense Artificial Intelligence Strategy*. Washington, DC: Department of Defense.
- [12] U.S.-China Economic and Security Review Commission (2026) *Two Loops: How China’s Open AI Strategy Reinforces Its Industrial Dominance*. Washington, DC: USCC.
- [13] RAND Corporation (2026) *U.S.-China Competition for Artificial Intelligence Markets*. Santa Monica, CA: RAND Corporation.
- [15] OECD (2025) *The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking*. Paris: OECD Publishing.
- [17] Pandise, E., Kent, J. L., Kaplan, M. and Mosk, M. (2026) ‘How the military is using AI in war’, *CBS News*, 18 March.
- [18] MIT Technology Review (2026) ‘Humans in the loop: the AI war illusion’, *MIT Technology Review*, 16 April.
- [20] Lieber Institute for Law and Warfare (2024) ‘AI-enabled targeting, Lavender, Gospel and decision-support systems under the law of armed conflict’, *Lieber Institute, West Point*.
- [21] Carnegie Endowment for International Peace (2026) ‘The fog of AI war’, *Carnegie Europe*, April.
- [22, 33] Amaral, N. (2026) ‘The Iran war highlights the creeping use of AI in warfare’, *Chatham House*, 27 March.
- [23] Stockholm International Peace Research Institute (2025) *Autonomous Weapon Systems and AI-Enabled Decision Support in Military Targeting*. Stockholm: SIPRI.
- [24] Stockholm International Peace Research Institute (2025) *Bias in Military Artificial Intelligence: Implica-*

tions for International Humanitarian Law. Stockholm: SIPRI.

- [25, 31] Human Rights Watch (2024) ‘Questions and Answers: Israeli Military’s Use of Digital Tools in Gaza’, *Human Rights Watch*, 10 September.
- [26, 32] Wirtschafter, V. (2026) ‘Generative AI as a weapon of war in Iran’, *Brookings Institution*, 8 April.
- [27] X (2021) *Community Notes*. X.
- [28] Nature (2026) ‘Stop the use of AI in warfare until laws are agreed’, *Nature*, March.
- [29] NATO (2024) *Revised NATO Artificial Intelligence Strategy*. Brussels: North Atlantic Treaty Organization.
- [34] Reuters (2026) ‘Drones dominate Ukraine battlefield four years into fighting’, *Reuters*, February.
- [35] Chomsky, M. (2026) ‘Armed Forces of Ukraine reports over 7,000 ground robotic system missions in January as frontline deployment expands’, *Defence Industry Europe*, 18 February.
- [36] Reuters (2026) ‘Ukraine opens battlefield data access to allies’ AI models’, *Reuters*, March.
- [37] Defense News Aerospace (2025) ‘Ukraine Emerges as World Leader in Drone Technology Driven by Battle-Proven Innovation’, 13 November.
- [38] Reuters (2026) ‘Romania tests AI-powered drone interceptors as Ukraine war gets closer’, *Reuters*, April.
- [39] Bajak, F. (2024) ‘US-China competition to field military drone swarms could fuel global arms race’, *Associated Press*, 12 April.