

INTENT AT HYPERSONIC VELOCITY: **NEW AND SHIFTING ETHICAL TRADE-OFFS IN AI-ENABLED AIR AND SPACE POWER**

DR. ALBERTO CHIERICI – NEW YORK UNIVERSITY, USA

THE VIEWS AND RESEARCH EXPRESSED IN THIS ARTICLE ARE MY OWN AND DO NOT REPRESENT THOSE OF MY EMPLOYER.

The integration of artificial intelligence (AI) into hypersonic defence systems presents unprecedented ethical and decision-making challenges. Compressed timelines, novel data signatures, and the convergence of air and space domains heighten pressures on human oversight, while opaque algorithms, latent data biases, and automated engagement decisions intensify accountability risks. Ethical frameworks, including Just War theory and AI safety principles, provide critical guidance for navigating these dilemmas. By embedding practical guardrails centred on transparency, accountability, and meaningful human control, this paper highlights how AI-enabled hypersonic capabilities can be harnessed responsibly, preserving both operational effectiveness and the moral legitimacy of air and space power in the decades to come.

Keywords

Artificial Intelligence, Hypersonic Weapons, Military Ethics, Autonomous Systems, Human–Machine Teaming, Accountability.

1. The Ethical Trade-offs of Autonomy at Hypersonic Speed

Hypersonic weapons and high-velocity platforms exemplify the speed-versus-oversight dilemma. These systems – such as boost-glide vehicles launched on suborbital trajectories – travel at speeds above Mach 5, manoeuvring unpredictably and drastically reducing the time for defenders to respond (Kunertova, 2024). In practical terms, a hypersonic missile can strike its target so quickly that human decision-makers have virtually no window to verify the threat or weigh collateral damage; by the time a person reacts, the engagement may be long over. This creates pressure to delegate detection, tracking, and even shoot/don't-shoot decisions to fast, autonomous algorithms, sometimes built using Artificial Intelligence (AI) techniques.

AI can be crucial for countering hypersonic threats: algorithms can fuse sensor data and coordinate interceptors far faster than human operators. AI-enabled early warning and air defence systems promise to overcome the limitations of human reaction time by instantly processing target trajectories and orchestrating responses in fractions of a second (Giovanardi et al., 2021). Automating lethal force decisions raises significant ethical concerns, even when it is against explicit military threats. These concerns pertain to human oversight, escalation risks, and accountability gaps, prioritising algorithmic

CONTACT: Dr. Alberto Chierici ✉ alberto.chierici@nyu.edu 📍 New York University, New York

Copyright Notice: © 2025 Capital 360 L.L.C. O.P.C. All rights reserved.

This academic journal is published by Capital 360 L.L.C. O.P.C. ("Capital 360") for the Dubai International Air Chiefs' Conference (DIACC) as a non-commercial, open-access publication available on the official event website for educational and research purposes. All content is the intellectual property of Capital 360 and/or the respective authors, with certain materials included under licence. Users may download, share, and cite the content for non-commercial purposes with proper attribution, provided that any reproduction, adaptation, or distribution for commercial use shall require the prior written consent of the applicable rights holder and/or Capital 360, as the case may be. This publication is governed by Federal Decree-Law No. 38 of 2021 on Copyright and Neighbouring Rights [UAE].

efficiency over nuanced human judgement and proportionality. The scaffolding of accountability – the deliberative process of verifying a target, checking proportionality, confirming orders – can collapse under hypersonic time pressures. Hypersonic velocity threatens to “outrun” our ability to apply moral and legal checks before a weapon is released. Without robust safeguards, the race for faster kill-chains could render the traditional norms of Just War (discrimination, proportionality, etc.) impractical artefacts of a slower era. The challenge is how to harness AI for decision speed without jettisoning human oversight – preserving a moral intent behind each action, even at machine tempo.

1.1. Opaque decision-making and diluted intentionality

AI systems increasingly exhibit what philosopher Luciano Floridi calls “agency without intelligence”. These algorithms can autonomously and effectively bring about particular outcomes without any understanding or intentionality behind the mathematical equations that govern their actions (Floridi, 2023).

In civilian contexts, we see AI models producing fluent language or identifying images with superhuman speed – yet they are essentially executing pattern recognition, not exercising judgement or comprehension (Bishop, 2021). Systems optimised for a machine objective often do not translate well into real-world objectives. For instance, well-known systems like the Generalised Pre-trained Transformers (GPTs, the technology behind tools such as ChatGPT) optimise the objectives of plausibility during pre-training and of working collaboratively with users during post-training. As a result, when used for different purposes, such as research and fact-finding, GPTs are known to hallucinate or become overly sycophantic (Hicks et al., 2024; Cheng et al., 2025). So, how can such systems be trusted in a military setting when the stakes are much higher? In the context of lethal air and space power, this separation of agency from intelligence raises non-trivial ethical dilemmas.

When a deep learning system powering a hypersonic missile autonomously identifies and targets an enemy satellite or aircraft, there is an outcome—the potential destruction of a high-value asset—but no moral actor in the loop truly intending that specific outcome. The algorithm has no concept of just cause or the value of human life; it cannot hold intent in the way a human decision-maker can. As a result, delegating lethal targeting to opaque AI dilutes intentionality in warfare.

The burden of intent and moral purpose, which traditionally rested with human commanders (“I chose to strike this target for these reasons”), becomes blurred when an AI’s inscrutable recommendation triggers the attack. Command responsibility may be legally assigned to the officers who deployed the system. But practically, it becomes harder to trace the line of volition – the chain of decision from human leadership to machine output. This raises concerns about an accountability gap.

If an autonomous drone makes the wrong decision (e.g., engages a civilian object misclassified as a threat), who is truly responsible? The commander did not intend that specific outcome, yet the machine that launched the strike has no capacity for intent or conscience. Military leaders pride themselves on

precision and discipline in the use of force; both rely on human judgement guiding each shot. An opaque AI, by contrast, may execute efficiently but without the ethical filters of guilt, doubt, or caution that temper a human’s trigger finger (Floridi, 2023).

Such “mindless agency” shifts the moral centre of military decision-making to an undefined place. Delegating decision-making to machines could be intentional to free the military from accountability, but also to avoid taking credit. In these scenarios, forces may transition from fighters to facilitators, which may require different personnel and training.

The Air Force’s professional ethos, which emphasises purposeful restraint (e.g., holding fire to avoid civilian harm or strategic error), is fundamentally challenged by systems that operate as black boxes. In short, because an AI cannot intend in any moral sense, using AI to conduct lethal operations risks actions without intent – a void of moral agency at the pointy end of the spear.

2. How Limitations of AI-Powered Automation Translate Into Recurring Military Hazards

Even with robust ethical principles and moral practice operationalised, militaries may face novel, recurring technical hazards when deploying AI in weapons and command systems. Three well-documented failure modes are particularly concerning: automation bias, latent data bias, and distribution-shift errors. These factors pose ethical trade-offs if not aggressively mitigated.

2.1. Automation Bias

Humans tend to over-trust automated systems, a phenomenon known as automation bias (Bode, 2025). In the context of hypersonic platforms utilising AI for threat assessment, this can mean operators deferring to an AI’s judgement without independent verification. For instance, if an AI misclassifies a non-hostile object as a high-priority, Mach 5+ hypersonic missile due to sensor anomalies or novel target signatures, an operator, influenced by automation bias, might approve a lethal response that they would otherwise have questioned. Evidence from civilian and military contexts shows people often default to automated suggestions. The 1988 downing of Iran Air Flight 655 by the USS *Vincennes*, whose crew misread radar data (Stewart, 2023), illustrates how automation-aided perception can cause fatal errors. This risk is significantly amplified with AI-driven hypersonic platforms due to their extreme velocity, which allows for even less time for human intervention and independent assessment. However, it’s worth noting that the potential for misidentification might be reduced with advanced autonomous systems designed to differentiate between distinct targets. The significant differences in data points between, for example, a civilian aircraft and a Mach 5+ hypersonic system could lessen the likelihood of such errors, potentially mitigating some concerns.

Countering automation bias requires training personnel to view AI as advisory, rather than authoritative, and designing interfaces that highlight uncertainty or offer alternatives. Maintaining “meaningful human control” means the operator must not become a passive confirmer of machine outputs, especially when

dealing with high-stakes, high-speed AI-driven platforms.

2.2. Latent Data Bias

AI systems learn from past data, which may contain hidden biases that reflect social or operational prejudices (Blanchard and Bruun, 2024).

In a military AI, bias in training data can result in skewed performance across various hypersonic platforms. Suppose the AI powering the guidance systems of a nuclear-capable hypersonic glide vehicle is trained on simulated flight data that disproportionately emphasises certain environmental conditions or threat profiles. In that case, it might develop a latent bias. This could lead to systematic degradation in accuracy or stability when encountering unrepresented conditions in real-world deployment, potentially causing misdirection or targeting errors. Similarly, in the development of hypersonic aircraft through AI-powered simulations by companies like Boeing, biases in the training data (e.g., incomplete or skewed historical flight data) could lead to machine learning models that inaccurately predict vehicle performance under certain critical scenarios, increasing the need for expensive and time-consuming physical tests. This latent bias in design and simulation data for hypersonic platforms could indirectly lead to operational vulnerabilities or an overreliance on human intervention to compensate for AI limitations.

Technical measures, such as dataset balancing, bias auditing, and algorithmic fairness techniques, are necessary in addressing this issue. Moreover, human operators must be aware of potential biases – e.g., an analyst should know if the AI guiding a hypersonic platform has higher error rates for specific categories of targets or environmental conditions, so they can compensate.

2.3. Distribution-Shift Failures

AI performance can deteriorate when it encounters conditions unlike its training data – a problem known as distribution shift (Ataei et al., 2021). This is particularly pertinent to hypersonic platforms. For example, AI guidance systems in hypersonic missiles might encounter unexpected atmospheric conditions or evasive manoeuvres not fully represented in their training data. Such errors could cause the missile to deviate from its intended trajectory or fail to reach its target. Similarly, autonomous hypersonic vehicles as Anduril's YFQ-44A Fury, while designed for high-G manoeuvres, could misinterpret novel sensor data from an unfamiliar combat environment, leading to unpredictable errors or misclassification of targets. The extreme speeds and unique flight profiles of hypersonic weapons inherently present novel data signatures that may not be adequately represented in training datasets. This could lead to misidentification or delayed responses from AI-powered defensive systems, thereby reducing their reliability.

The reported use of an autonomous drone in North Africa in 2020—allegedly engaging a target without direct human authorisation—underscores the risks associated with autonomous systems operating in

unforeseen circumstances, a risk amplified by the compressed decision cycles inherent to hypersonic warfare (Wehrey and Bonney, 2025).

To mitigate distribution shift issues, AI models for hypersonic platforms should undergo extensive stress testing and “red-teaming”: exposing them to varied, worst-case scenarios and adversarial inputs to identify where they fail.

For instance, simulation frameworks developed by major aerospace firms such as Boeing could be expanded to encompass a wider range of extreme and diverse environmental conditions than those currently modelled. Similarly, an AI system employed for data processing on platforms like the Triton high-altitude surveillance drone may encounter inputs that deviate significantly from its training parameters—hence the importance of programming such systems to recognise statistically novel scenarios.

In that case, it should raise a flag or defer to human control rather than proceeding with high confidence. No AI should be trusted in a domain until it’s proven across the full range of conditions likely to be faced or its deployment is sufficiently limited to address potential risks – and even then, a human commander must be ready to intervene when the unexpected arises. Human judgement must remain a fallback whenever confidence is low.

Beyond these failure modes, other technical concerns include cybersecurity risks (adversaries hacking or spoofing AI systems) and unpredictable AI interactions (two opposing AI systems might escalate a conflict in unforeseen ways). For instance, an autonomous air defence AI might misinterpret a rival’s manoeuvres and trigger a retaliation that spirals – an “algorithmic escalation” scenario. These are not hypotheticals: a 2024 review of AI in hypersonic defence noted that adversaries could exploit vulnerabilities in AI, and autonomous systems might inadvertently escalate conflicts due to misinterpretations of threats if not carefully designed (Zohuri2024). Such challenges necessitate rigorous testing, ethical oversight, and robust fail-safes in any AI-enabled military platform.

3. A Case Study on AI in Military Operations

Exercises simulating high-end conflict have shown the significant impact of AI-enabled decision aids.

Human–Machine Teaming in the Air Force. The U.S. Air Force has trialled Advanced Battle Management System (ABMS) and related AI battle-management tools through a series of exercises, including the “Decision Advantage Sprints for Human–Machine Teaming” (DASH) wargames. In the first DASH exercise, officers from allied countries simulated a high-end conflict twice – once using traditional methods and once with AI decision aids integrated into their command-and-control setup (Freedberg 2025). An Air Force-developed AI, known as the Transformational Model for Decision Advantage, suggested courses of action, including target sequences, asset retasking, and logistics moves.

With AI assistance, the staff addressed twice as many operational dilemmas and produced three times as many viable solutions within the same timeframe. Overall throughput reportedly increased sevenfold, without any decline in quality. In other words, the human–AI team made faster decisions with accuracy comparable to that of human-only teams.

Concerns about AI hallucinations or spurious errors did not significantly materialise; error rates were “on par with human error” according to Air Force analysis. The trial highlighted AI’s potential to accelerate decision-making in the kill chain without undermining effectiveness.

Crucially, humans remained in the loop – the AI proposed options, but officers still validated and selected them. Col. Christopher Cannon, head of the ABMS team, emphasised that the goal was not to replace humans but to actively assist operators in transforming data into more informed battle management decisions. This reflects a doctrinal commitment to centaur-like teaming (human + AI) rather than full autonomy in command and control.

Nonetheless, even these controlled successes highlight issues for the future. One issue is operator trust: early on, if an AI suggestion seemed counterintuitive, some officers hesitated – a reminder that trust calibration (neither blind faith nor distrust) is key. Another issue is how to scale these tools securely; as more data and decisions get handled by AI, the system’s cybersecurity and integrity become mission-critical (you wouldn’t want an adversary hacking your battle management AI to feed false recommendations). Finally, validation in peacetime exercises is one thing; performance under adversarial pressure, electronic warfare, and partial information fog is another. These experiments offer a valuable sandbox for understanding the power of accelerated AI decision-making and the necessity of robust guardrails before wartime use particularly as nations pursue similar efforts for advanced military capabilities.

4. Bridging Just War Principles and Operational Practice

To navigate the challenges explored thus far, classical Just War principles – just cause, right intention, discrimination, and proportionality – must be translated into the concrete language of air operations and AI system design. This means ensuring that age-old ethical criteria can be built into the rules of engagement, algorithms, and workflows that guide AI-enabled combat.

Just Cause & Right Intention. At the strategic level, AI should only be employed in conflicts and missions that serve a legitimate defensive purpose (just cause) and with the proper intent (e.g., protecting peace and security, rather than aggression or reprisal). Practically, this translates to tight human control over when autonomous systems are activated to use force. A human commander must still decide why an engagement is initiated and ensure it aligns with lawful objectives – an AI should never be left to decide whether to attack purely on the basis of algorithmic logic. For example, a target recommendation system might flag potential threats in a hypersonic context. Still, the decision to prosecute those targets must be driven by a human-validated mission rationale (e.g., neutralising an imminent military threat),

not by the AI's own optimisation goals. Maintaining a human-in-the-loop for the ultimate initiation of lethal force, supported by explainability features, robust audit trails, and rigorous red-teaming exercises (as discussed in Section 5), helps preserve the right intention – the strike occurs as a conscious act of policy and duty, not as a side effect of a software process.

Distinction. Perhaps the most critical operational principle is distinguishing combatants from non-combatants and lawful targets from protected people/objects. AI can aid distinction by improving target recognition – indeed, the U.S. military's Project Maven “made extensive use of algorithms” in Iraq and Syria to classify objects in drone imagery, e.g., tanks vs. trucks vs. civilian vehicles (Wehrey and Bonney, 2025). Such tools, if properly trained, can enhance precision and reduce accidental strikes. In less cluttered domains or domains where threats have very distinct profiles – e.g., responding to hypersonic threats – distinction may also be more easily satisfied by machines or AI-enabled systems. However, distinction is only as good as the data and rules the AI is given. As we pointed out earlier, biased or insufficient training data can cause an AI to misidentify civilians as combatants, with deadly consequences (Bode, 2025). Commanders and developers shall ensure that distinction criteria are encoded in AI systems and rigorously tested. This could involve embedding rules that specify categories of objects (e.g., ambulances, schools, hospitals) marked with proper symbols that are off-limits – essentially, no-strike lists hardwired into the AI's targeting filters. It also means keeping a human validator in or on the loop to review AI-generated target selections, especially in complex environments. The concept of “positive identification (PID)” – requiring high confidence that a target is legitimate before engagement – remains a cornerstone. AI can assist in PID by rapidly cross-referencing sensor data, but a human should confirm that identification unless in extreme point-defence scenarios. The broad strategic implications of hypersonic threats, including their impact on deterrence, offence, and the balance of power, necessitate a re-evaluation of national security strategies beyond individual asset protection. The ethical mandate to never deliberately target innocents must be translated into both software (through robust classification models and constrained algorithms) and procedure (through human oversight of AI target recommendations).

Proportionality. This principle requires that the harm caused by any military action not be excessive relative to the concrete and direct military advantage gained. In operational terms, even a lawful target should only be engaged in a way that minimises collateral damage. AI can help here by running faster collateral damage estimates and suggesting munitions that yield the desired effect with the least risk to civilians. For example, an AI might evaluate multiple strike options (different weapon types, angles of attack, timing) and highlight the one that achieves the objective with the lowest estimated civilian impact. However, proportionality is also a deeply contextual and value-driven judgment that an AI cannot make on its own – it involves weighing intangibles like the risk to civilian life against the urgency of neutralising a threat. The concern is that an algorithm might treat this as a straightforward optimisation (e.g., a formula balancing casualties vs. target value) without grasping the moral weight behind those variables. To keep proportionality grounded in human values, militaries can incorporate human-in-the-loop approval for strikes above certain risk thresholds. For instance, if an AI-driven battle management system projects that striking a target could cause significant collateral damage, it should

flag this for the commander's review rather than automatically engaging.

Additionally, AI systems should be “tuned” to err on the side of caution – a bias against action in ambiguous cases – reflecting the idea that when in doubt, one should refrain from using lethal force. In domains like hypersonic threat interception, proportionality will often not present complex assessment needs, as the interception of missiles at high altitude rarely involves direct collateral harm. However, risks imposed by falling debris will demand accounting in proportionality, centring the need for nuanced judgment regarding these potential harms to protected persons and objects.

The timeless tenets of just-war doctrine remain applicable in the age of AI and hypersonic weapons – but they must be integrated into the algorithms and standard operating procedures (SOPs) of modern warfare. Air chiefs and planners should treat ethical criteria as equally mission-critical to technical criteria. For example, an AI-enabled kill chain should not only be evaluated on its speed and lethality, but also on whether it reliably upholds discrimination and proportionality standards in its recommendations.

This also highlights the urgent need for legislative norms and doctrinal support to evolve in tandem with technological advancements. If militaries increasingly shift from fighters to facilitators – as previously noted – then laws, training standards, and doctrine must be updated to clarify accountability, codify ethical requirements, and prevent gaps in responsibility (cfr Woods, 2023). Embedding such norms into both international law and national doctrine would help ensure that rapid advances in autonomy do not outpace the frameworks that bind force to justice.

5. Technical and Policy Guardrails for Accountable AI

How can AI-enabled warfare remain accountable? Several technical and policy guardrails are being pursued or recommended within defence departments (Air Force, 2025):

Testing and Validation: Should AI face the same certification as aircraft or weapons? Stress tests, red-teaming, and accreditation schemes can reduce fragility and adversarial risk. Regular re-validation is vital, especially with technology that evolves fast and adapts to new data. Technically, this means developing advanced simulation environments that account for the extreme speeds, complex atmospheric interactions, and potential for rapid target re-prioritisation inherent in hypersonic systems. Accreditation schemes should include ethical review boards and independent evaluators to certify that AI systems meet predefined standards for target discrimination and adherence to the laws of armed conflict.

Audit Trails: Every AI decision should leave a traceable record: what the AI saw, what it recommended, and what humans decided. Would commanders act differently if they knew their actions were being logged and were subject to review?

For a hypersonic context, this includes real-time telemetry, target acquisition data, and the specific algorithms and parameters used for decision-making. The ethical impact of knowing actions are logged is profound: it encourages commanders to act with greater deliberation, knowing their decisions will be subject to retrospective ethical and legal review. This data is critical for post-conflict analysis, allowing

for the identification of potential biases, errors, or violations of international law.

Explainability: AI tools should provide reasons or confidence levels for their outputs. Can operators see why a system flagged a target, and challenge it if needed? Documentation of training data, authorisation, and failure modes should be standard. Wood (2024) presents a comprehensive argument regarding the role of explainability within a military context. In a hypersonic scenario, this could involve the AI explaining why it prioritised one target over another, or why it identified a specific object as a legitimate military objective given the available data. Documentation of training data, authorisation processes, and identified failure modes enable operators to identify and challenge potentially erroneous or biased AI decisions.

Governance and Human Override: Good governance requires robust rules of engagement that clearly delineate when AI can execute actions independently (e.g., trajectory adjustments to maintain speed) and when human approval is non-negotiable (e.g., final weapon release). Fail-safe designs and easily accessible “off switches” are critical, providing an immediate means to halt AI operations if unforeseen circumstances or ethical concerns arise.

The ethical trade-off lies in balancing the speed requirements of hypersonic warfare with the need for deliberate human decision-making. Forward-looking doctrine must establish clear lines of authority and responsibility for AI-enabled systems, ensuring that human commanders retain ultimate ethical accountability and responsibility. Governance structures will need to develop international norms and treaties to regulate the development and deployment of autonomous hypersonic weapons, potentially including outright prohibitions on fully autonomous lethal systems. Standards will need to be developed for the design and implementation of fail-safe mechanisms and human-in-the-loop protocols.

Training and Ethics: The ethical foundation of secure AI deployment rests on well-trained and ethically informed human judgment. No technical guardrail can entirely replace the human capacity for moral reasoning and adaptation. Training programs must go beyond technical proficiency, actively cultivating critical thinking skills to question AI outputs, recognise potential biases (e.g., in target identification), and understand the ethical implications of AI-driven decisions in high-stakes, time-compressed scenarios. For hypersonic operations, this means training personnel to understand the unique challenges of attribution, de-escalation, and proportionality when dealing with speedy and potentially evasive threats. The ethical goal is to empower personnel to exercise moral courage and take control when AI systems demonstrate behaviour that deviates from ethical norms or international law.

These measures are not exhaustive, but together they point toward a practical framework: AI systems that are transparent, tested, logged, governable, and used by ethically trained humans. The open question is whether militaries will implement them rigorously before conflict forces the issue.

6. Conclusion

The integration of AI into hypersonic defence systems raises enduring ethical questions in a new technological setting. Opaque algorithms, biased data, and accelerated decision cycles together create accountability gaps and weaken the human role in decision-making. While such challenges are not unique to hypersonics, the speed and complexity of these systems intensify the risks and demand more precise ethical boundaries.

Principles of transparency, accountability, and meaningful human control must be preserved as essential guardrails. Looking forward, as air and space domains converge and hypersonic systems evolve, future research should explore how governance, verification mechanisms, and technical standards can be developed internationally. Ensuring ethical AI in hypersonic contexts will not only shape operational outcomes but also the credibility and legitimacy of air and space power itself.

References

- Air Force (2025). Air Force Doctrine Note 25-1, Artificial Intelligence. [online] Available at: https://www.doctrine.af.mil/Portals/61/documents/AFDN_25-1/AFDN%2025-1%20Artificial%20Intelligence.pdf. Accessed on: 29 August 2025.
- Ataei, M., Erdogdu, M., Kocak, S., Ben-David, S., Saleh, S., Ghazi, A., Nguyen, J., Khayrat, K., Pesaraghader, A., Alberts-Scherer, A., Sanchez, G., Pouryazdian, S. and Zhao, B. (2021). Understanding Dataset Shift and Potential Remedies. A Vector Institute Industry Collaborative Project Technical Report. Vector Institute. [online] Available at: https://vectorinstitute.ai/wp-content/uploads/2021/08/ds_project_report_final_august9.pdf. Accessed on: 29 August 2025.
- Bishop, J.M. (2021). Artificial Intelligence Is Stupid and Causal Reasoning Will Not Fix It. *Frontiers in Psychology*, 11. DOI: <https://doi.org/10.3389/fpsyg.2020.513474>.
- Blanchard, A., and Bruun, L. (2024). Bias in Military Artificial Intelligence. Stockholm International Peace Research Institute Background Paper. SIPRI. [online] Available at: https://www.sipri.org/sites/default/files/2024-12/background_paper_bias_in_military_ai_0.pdf. Accessed on: 29 August 2025.
- Bode, I. (2024). The Problem of Algorithmic Bias and Military Applications of AI. [online] *Humanitarian Law & Policy Blog*. Available at: <https://blogs.icrc.org/law-and-policy/2024/03/14/falling-under-the-radar-the-problem-of-algorithmic-bias-and-military-applications-of-ai>. Accessed on: 29 August 2025.
- Carchidi, V., and Soliman, M. (2023). The technical is geopolitical: Expanding US-UAE relations through AI. [online] Middle East Institute. Available at: <https://mei.edu/publications/technical-geopolitical-expanding-us-uae-relations-through-ai>. Accessed on: 29 August 2025.
- Cheng, M., Yu, S., Lee, C., Khadpe, P., Ibrahim, L. and Jurafsky, D., 2025. Social Sycophancy: A Broader Understanding of LLM Sycophancy. arXiv preprint arXiv:2505.13995.
- Floridi, L., 2023. AI as Agency Without Intelligence: on ChatGPT, Large Language Models, and Other Generative Models. *Philosophy & technology*, 36(1), p.15. DOI:<https://doi.org/10.1007/s13347-023-00621-y>.
- Freedberg, S.J. (2025). AI Transformational Model accelerates battle staff decision-making ‘seven-fold’ in Air Force experiment. [online] *Breaking Defense*. Available at: <https://breakingdefense.com/2025/06/ai-transformational-model-accelerates-battle-staff-decision-making-seven-fold-in-air-force-experiment/>. Accessed on: 29 August 2025.

Giovanardi, M., Trane, M. and Pollo, R., 2021. IoT in Building Process: A Literature Review. *Journal of Civil Engineering and Architecture*, 15(9), pp.475-487.

Hicks, M.T., Humphries, J. and Slater, J., 2024. ChatGPT is bullshit. *Ethics and Information Technology*, 26(2), pp.1-10. DOI: <https://doi.org/10.1007/s10676-024-09775-5>.

Kunertova, D., 2021. Hypersonic Weapons: Fast, Furious... and Futile?. *RUSI Newsbrief*, 41(8). Available on <https://www.rusi.org/explore-our-research/publications/rusi-newsbrief/hypersonic-weapons-fast-furiousand-futile>. Accessed on 29 August 2025.

Stewart, B.S. (2023). USS Vincennes Shoots Down Iranian Civilian Plane. EBSCO Information Services, Inc. [online] Available at: <https://www.ebsco.com/research-starters/military-history-and-science/uss-vincennes-shoots-down-iranian-civilian-plane>. Accessed on: 29 August 2025.

Wehrey, F., and Bonney, A. (2025). The Middle East's AI Warfare Laboratory. *War on the Rocks*. [online] Available at: <https://warontherocks.com/2025/04/the-middle-east-s-ai-warfare-laboratory>. Accessed on: 29 August 2025.

Wood, N.G., 2023. Autonomous weapon systems and responsibility gaps: a taxonomy. *Ethics and Information Technology*, 25(1), p.16. DOI: <https://doi.org/10.1007/s10676-023-09690-1>.

Wood, N.G., 2024. Explainable AI in the military domain. *Ethics and Information Technology*, 26(2), p.29. DOI: <https://doi.org/10.1007/s10676-024-09762-w>.

Zohuri, B. (2024). Harnessing Artificial Intelligence for Countering Hypersonic Weapons: A New Frontier in Battlefield Offense and Defense (A Short Review). *Journal of Energy and Power Engineering*, 18(4). DOI: <https://doi.org/10.17265/1934-8975/2024.04.002>.