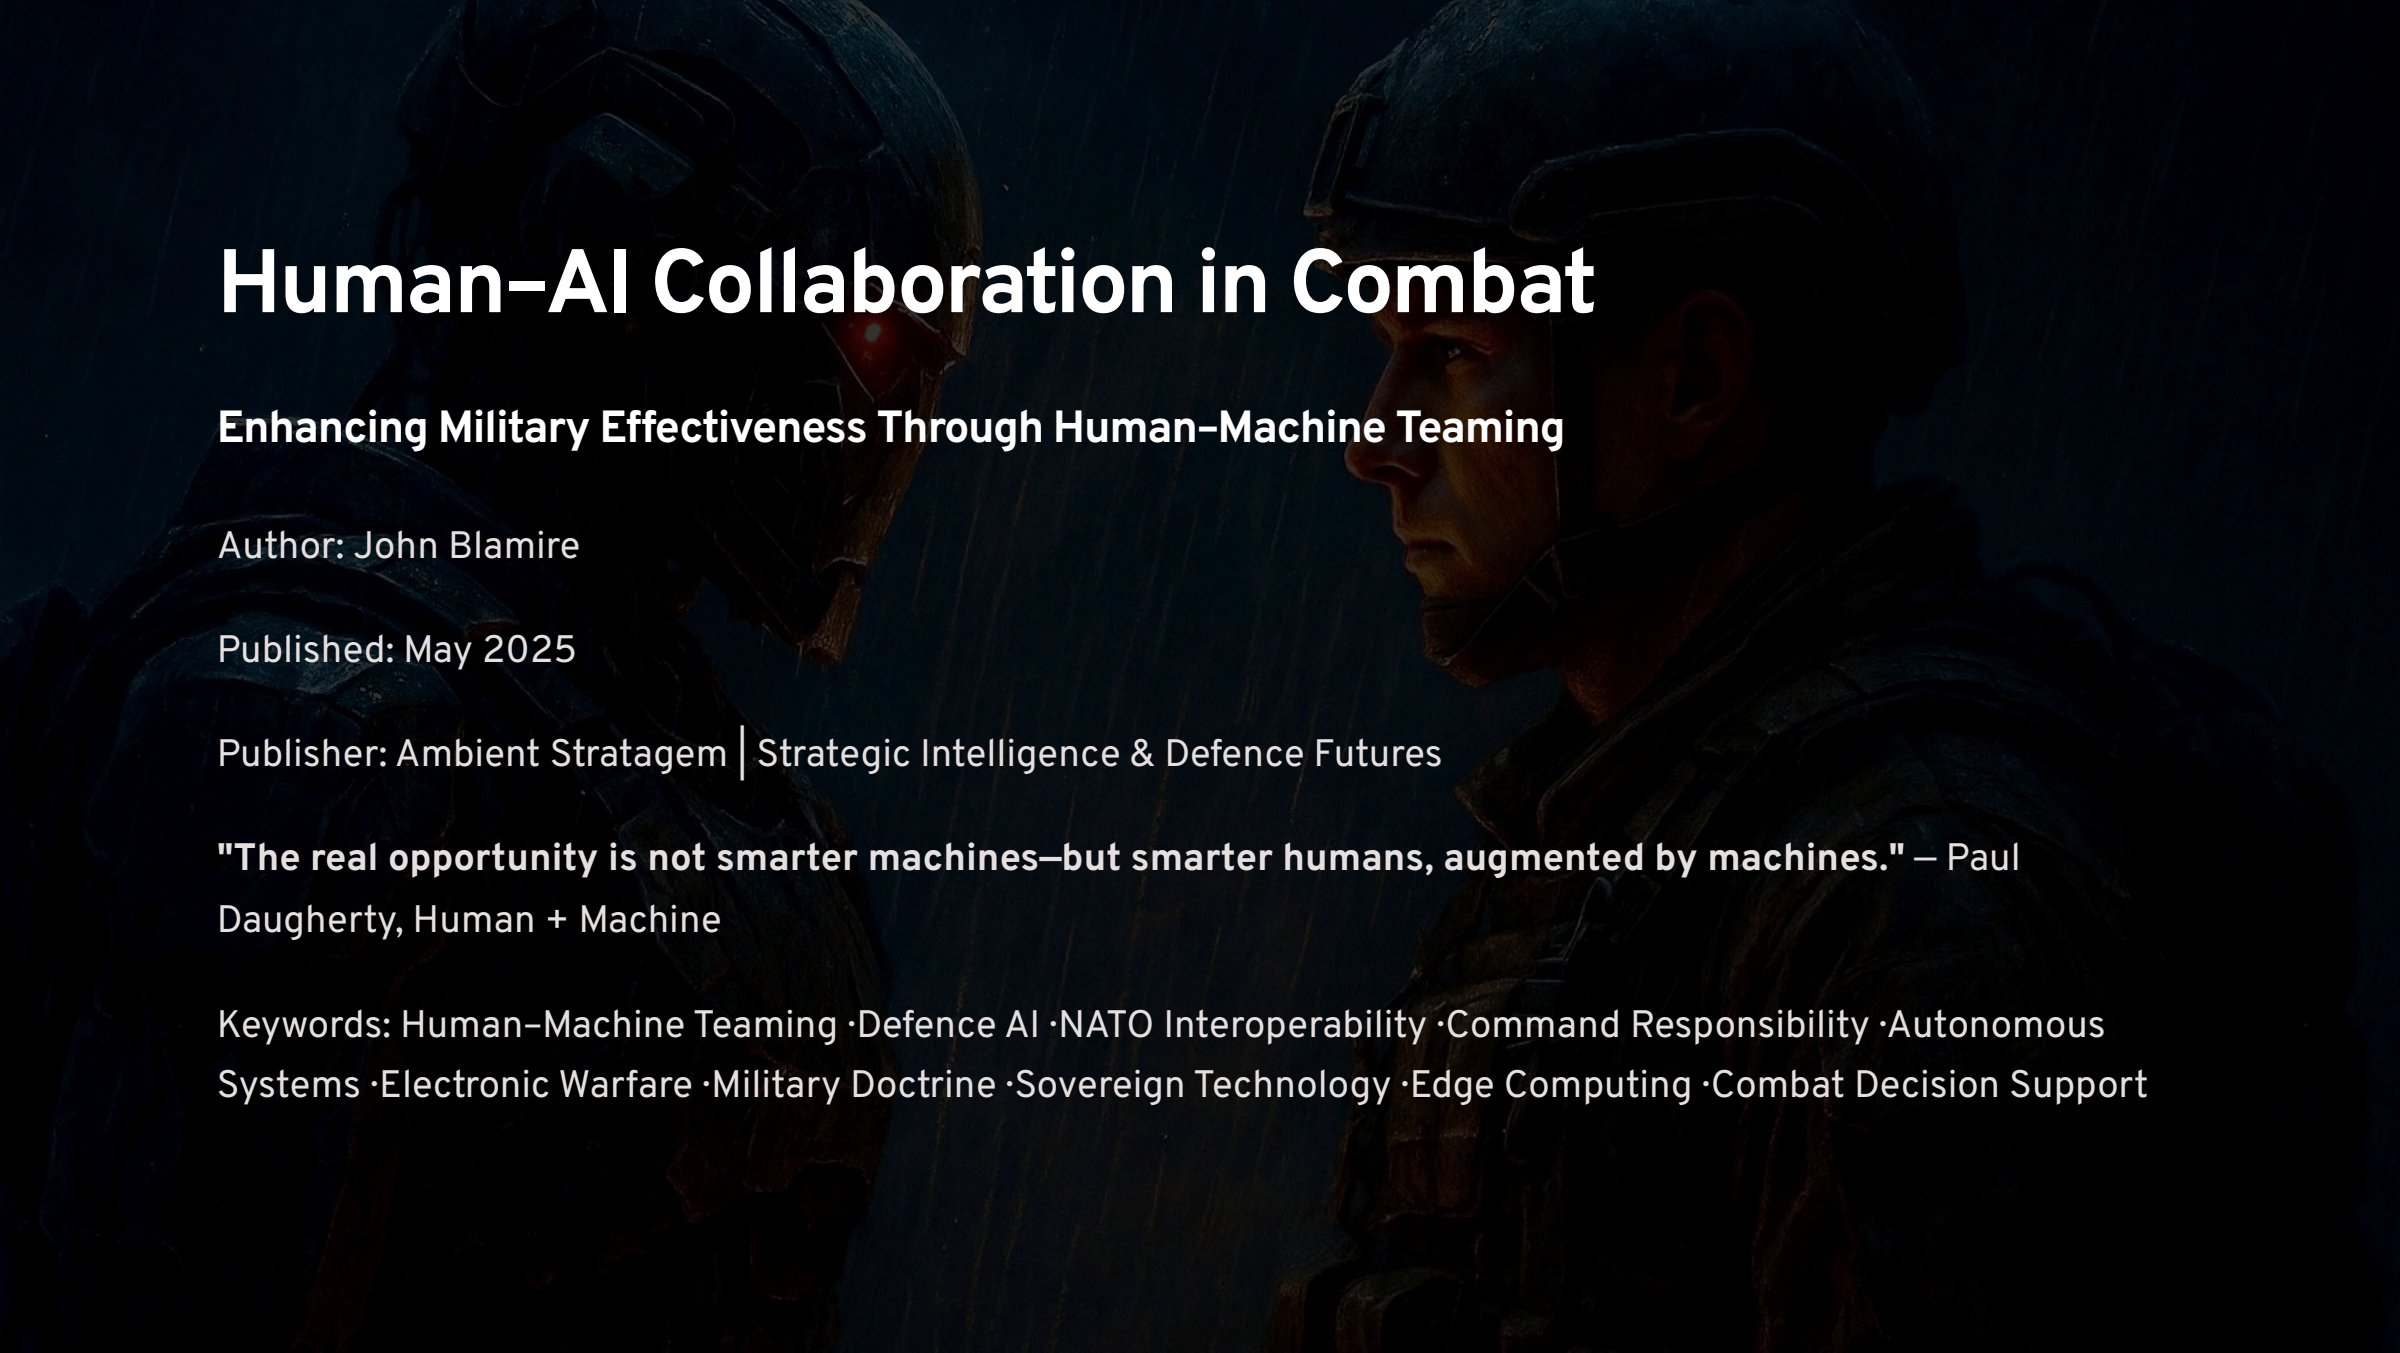

Human–AI Collaboration in Combat

Enhancing Military Effectiveness Through Human–Machine Teaming

Author: John Blamire

Published: May 2025

Publisher: Ambient Stratagem | Strategic Intelligence & Defence Futures

"The real opportunity is not smarter machines—but smarter humans, augmented by machines." — Paul Daugherty, Human + Machine

Keywords: Human–Machine Teaming ·Defence AI ·NATO Interoperability ·Command Responsibility ·Autonomous Systems ·Electronic Warfare ·Military Doctrine ·Sovereign Technology ·Edge Computing ·Combat Decision Support

Executive Summary

In an era marked by mounting global instability, declining confidence in traditional alliances, and the return of great power confrontation, military advantage hinges not merely on hardware or personnel—but on the fusion of both through intelligent, adaptive systems. This white paper explores the strategic imperative for Human–AI Collaboration in Combat, offering a grounded yet forward-looking assessment of how human-machine teaming can enhance operational effectiveness, lethality, and survivability across all domains of warfare.

As the US commitment to European defence becomes increasingly contingent, and NATO cohesion faces internal and external pressures, sovereign defence capabilities must evolve rapidly. The battlefield of the future will not be won by autonomous machines acting in isolation, nor by unaided human decision-making. It will be shaped by human operators empowered by embedded AI, capable of making faster, better-informed decisions even in communications-degraded or electronic warfare (EW)-contested environments.

Drawing on lessons from programs such as the U.S. Department of Defense's Project Maven, as well as insights from the RAND Corporation and leading technologists like Paul Daugherty, this white paper argues that ambient AI—not autonomy—is the key enabler of next-generation military advantage.

> "Machines are for answers. Humans are for questions." — Kevin Kelly, Founding Executive Editor of Wired

Key Takeaways

- Human-machine teaming is a force multiplier—reducing cognitive burden while accelerating battlefield response time.
- Sovereign, trustworthy AI systems must be fielded at the edge, with humans firmly in the loop, especially in EW-denied conditions.
- The UK and European defence ecosystems must act decisively to avoid technological dependence on external powers or commercial monopolies.
- A paradigm shift is required: from 'autonomous systems' to human-embedded AI advisors, trained on domain-specific tactical logic.

This white paper is designed for:

- Defence agencies seeking to operationalise AI without ceding control.
- Aerospace and defence contractors pursuing next-generation mission systems.
- Strategic thinkers shaping doctrine, procurement, and force design for 2030 and beyond.

The Tactical Reality – Why Human–AI Teaming is No Longer Optional

Across the battlefields of Ukraine, the deserts of the Sahel, and the Indo-Pacific's contested sea lanes, a new truth has emerged: traditional command-and-control structures are being overwhelmed. The speed, complexity, and volume of modern warfare now routinely exceed human cognitive limits. To prevail in this environment, armed forces must not only digitise—they must team.

From Superiority to Survivability

The modern battlefield is no longer linear. It is multi-domain, constantly shifting, and increasingly saturated with data. The adversary—be it state or non-state—often adapts faster than our institutions can respond. In such contexts, survivability itself now depends on decision advantage: the ability to sense, decide, and act faster and more accurately than the enemy.

- "In war, the moral is to the physical as three to one."
- Napoleon Bonaparte

Today, the informational is to the kinetic as three to one. The side that better integrates information into actionable decisions—in real time—gains the edge.

Human–AI collaboration is no longer a futuristic vision. It is the only viable path to achieving this edge in contested environments where:

- Communications are intermittent or denied.
- Electronic warfare (EW) degrades GPS, cloud, and satcom reliance.
- Tactical decisions must be made under fire, with seconds to spare.

Human Overmatch in a Machine-Speed Fight

Despite advances in autonomy, the core military challenge remains unchanged: deciding who does what, when, and why, under uncertainty. This is where machines alone still falter.

AI excels at pattern recognition, probability, and computation at speed. But it lacks intent, context, and value judgement—elements that remain uniquely human. The future of warfare lies in fused cognition, where each side of the human–machine partnership contributes its unique strengths:

- Humans: ethical judgement, mission intent, adaptive improvisation.
- AI: data fusion, threat recognition, predictive analysis, precision execution.

The U.S. Army's Project Maven offered an early glimpse into this dynamic. By integrating computer vision models to support ISR analysts in Afghanistan, it reduced video analysis time from hours to seconds—yet always left the decisional authority in human hands.

The European Imperative

In Europe, the need for this approach is particularly acute. Faced with:

- An overstretched US, whose attention is pivoting to the Pacific.
- A revanchist Russia, probing NATO's resolve and exploiting digital seams.
- Rapid Chinese dual-use tech diffusion, especially in drones and comms.

European states must now prioritise the development and fielding of sovereign AI-enabled combat systems that reinforce human agency rather than replace it.

This requires rejecting the Silicon Valley narrative of full autonomy and instead embracing a warfighting AI model rooted in trust, operational nuance, and frontline empowerment.

Building Trust into the Machine – Ethics, Accountability, and Command Responsibility

Trust is the decisive factor in battlefield AI adoption. It is not enough for a system to work in the lab—it must be trusted in the field, under fire, and by commanders responsible for lives and missions. This chapter explores how trust is built into human–AI teaming architectures, not just through performance, but through transparency, accountability, and ethical alignment with military command responsibility.

The Weight of Command

Unlike commercial contexts, military leaders operate under laws of armed conflict, rules of engagement, and deep moral accountability. Every decision carries operational, legal, and political consequences. Delegating meaningful control to a machine without a robust understanding of how it reasons, or under what constraints, is a risk few commanders will take.

> "Responsibility is indivisible. If a machine errs, the human commander bears the burden."

— Adapted from Admiral James Stavridis

This reality means that for human–AI collaboration to succeed, systems must be designed with trust as a feature—not an afterthought.

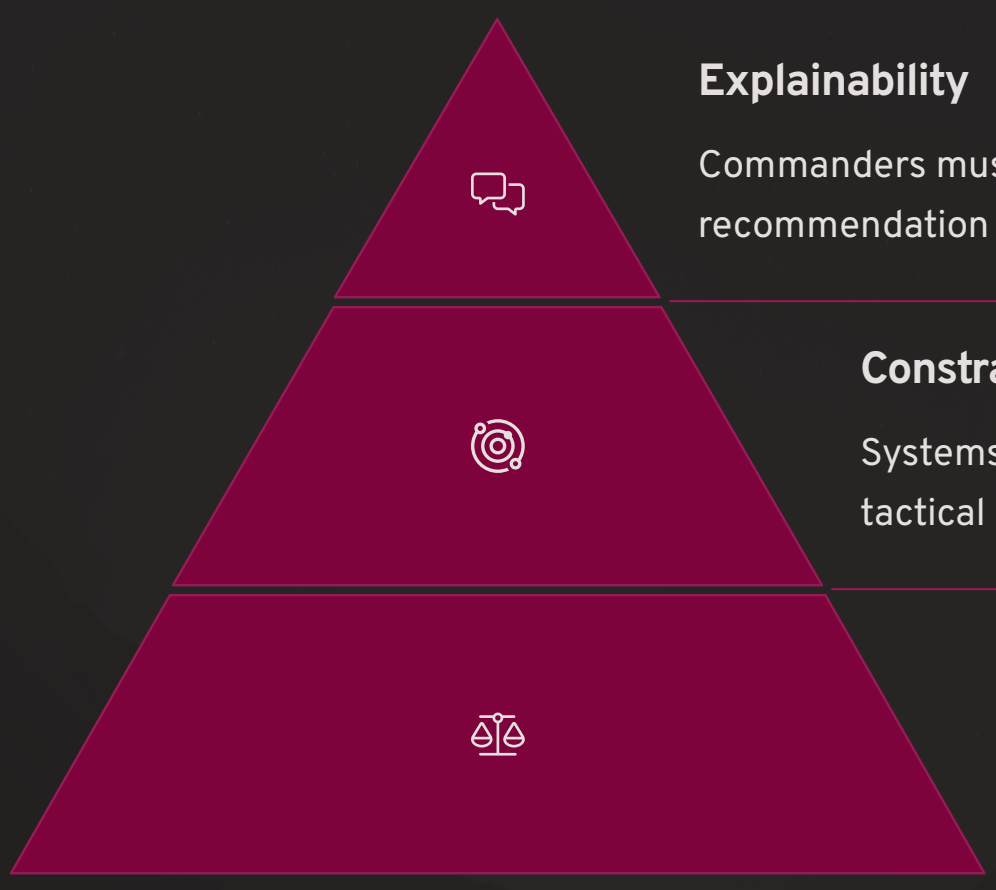

Explainability

Commanders must understand how a system reached a recommendation or decision

Constraint

Systems must operate within defined mission parameters and tactical rulesets

Accountability

Every AI action must leave a traceable record

In short: AI must advise, not decide—and must always operate as a transparent, bounded partner within the chain of command.

Lessons from NATO and Allies

Some NATO states have already begun grappling with these issues:

- The UK's Defence AI Strategy explicitly commits to human-centred autonomy, rejecting lethal autonomous weapons that operate without meaningful human control.
- France's doctrine focuses on le combattant augmenté—the augmented soldier—positioning AI as a force enhancer, not a substitute.
- Germany's Bundestag has called for legally binding limits on AI weapons, particularly in relation to accountability and ethics.

However, few have yet operationalised these doctrines into deployable, trusted edge systems. This gap presents both a vulnerability and an opportunity for leadership.

Avoiding the Autonomy Trap

The current discourse is skewed by two unhelpful extremes:

- On one end, utopian technologists pushing for fully autonomous kill chains, often detached from military reality.
- On the other, risk-averse institutions paralysed by ethical overcorrection, treating all AI as untrustworthy.

The strategic middle ground—trusted AI teaming embedded within human judgement structures—remains underdeveloped, yet it is here that operational advantage lies.

Logic Under Fire – The Technical Demands of Combat-Grade AI

Not all AI is created equal. The systems powering online ads, consumer voice assistants, or commercial logistics operate under vastly different conditions than those expected in warfare. Combat-grade AI must function amidst chaos, deception, degradation—and it must do so with absolute reliability.

This chapter sets out the core technical requirements that distinguish battlefield-ready AI from conventional enterprise systems. At its heart lies one principle: mission logic must survive contact with the enemy.

The Nature of the Combat Environment

Modern military operations are defined by contested, unstable, and unpredictable environments. Combat AI must function under:

- Electronic warfare (EW): GPS jamming, comms interference, signal denial.
- Intermittent connectivity: Networks are degraded, denied, or compromised.
- Hostile physical conditions: Dust, heat, vibration, mobility, and destruction.
- Cognitive overload: Humans under stress, fatigue, and constant decision pressure.

The implication is clear: AI must be embedded, resilient, and context-aware, rather than cloud-dependent, fragile, or over-generalised.

> "The enemy gets a vote. Any AI that ignores this reality belongs in a lab, not a warzone."

— General Stanley McChrystal (adapted)

Key Technical Demands

Edge Execution

Combat AI must run at the edge—on lightweight, power-efficient, ruggedised hardware within tactical vehicles, UAVs, or soldier systems. It cannot rely on cloud infrastructure or constant connectivity.

Low Latency, High Confidence

In combat, decisions often have to be made in sub-seconds. The AI system must provide high-confidence outputs fast—without 'spinning' or failing silently.

Determinism and Verifiability

Systems must behave predictably under stress. Deterministic, logic-based systems—like stateful dataflow graphs—offer traceable, testable decision paths, which are essential in a military context where outcomes must be understood and justified.

Mission-Specific Adaptability

Combat AI should not attempt to be general-purpose. It must be trained and tuned on specific mission profiles, terrain models, enemy TTPs (tactics, techniques, procedures), and force doctrine. Adaptability is key—but bounded by mission logic.

Offline Learning and Updating

Since battlefield systems may be offline for extended periods, AI updates must be portable, verifiable, and securely propagatable via field devices or cross-domain data links.

Why Most AI Systems Fail in Combat

The majority of commercial AI systems—especially LLMs and foundation models—fail under military constraints due to:

- Bandwidth dependency: Requiring constant cloud access.
- Opacity: Inability to explain or audit decisions.
- Generalisation: Trained on the internet, not on adversary tactics.
- Vulnerability: Susceptibility to spoofing, noise, and adversarial inputs.

This is why a new generation of AI architectures is required—purpose-built for resilience, logic integration, and tactical relevance.

From Augmented Soldier to Intelligent Force – Scaling Human–AI Teaming Across the Battlespace

The promise of Human–AI collaboration does not lie in isolated use cases. Its true strategic value emerges only when scaled across units, formations, and operational theatres—creating an intelligent, adaptive force where every echelon benefits from real-time decision advantage.

This chapter explores how AI-enhanced logic can move from the individual warfighter to the collective battlespace, enabling what might be called cognitive manoeuvre—the ability to out-think, out-react, and out-adapt the adversary at every level of conflict.

> "To subdue the enemy without fighting is the acme of skill."

— Sun Tzu

In the 21st century, this skill hinges on out-decisioning the enemy at machine speed, with human intent at its core.

The Augmented Soldier

Individual combatant with AI assistance for threat detection, navigation, and tactical decisions

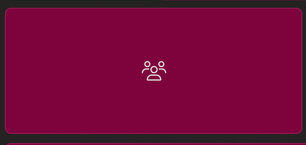

Distributed Teaming

Sections and squadrons sharing logic and maintaining tactical alignment

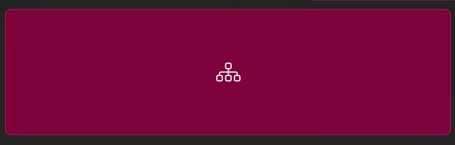

Cognitive Mesh

Battlegroups and brigades with collective intelligence and adaptive capabilities

The Augmented Soldier

At the tactical edge, the human–AI partnership begins with the individual combatant. This is not science fiction. In practical terms, AI can already assist soldiers with:

- Real-time threat detection through sensor fusion and contextual alerting.
- Navigation and terrain analysis in GPS-denied environments.
- Combat ID and fratricide prevention via visual and acoustic tagging.
- Tactical decision support through mission-specific rule advisories.

These capabilities are not meant to automate combat but to unburden the operator, sharpening attention and increasing survivability.

> "What AI gives the warfighter is not omniscience—but clarity, when it counts."

— Defence AI Architect, 2023 NATO Symposium

From Section to Squadron: Distributed Teaming

Scaling from the soldier to the section or squadron requires AI systems to:

- Share logic, not just data—allowing common decision frameworks across units.
- Operate with consistency under fire, even as individual nodes are degraded or destroyed.
- Support command intent, ensuring decentralised actors remain tactically aligned.

By embedding tactical logic in a machine-readable, portable format, units can maintain coherence even when isolated—a key requirement in peer-to-peer conflict where comms cannot be guaranteed.

Battlegroup and Brigade: A Cognitive Mesh

At higher echelons, human–AI teaming evolves from individual augmentation to collective intelligence. Key enablers include:

- Edge-federated learning: Tactical nodes sharing battlefield insights asynchronously, enabling adaptation without cloud dependency.
- Mission logic propagation: Distributing refined SOPs (standard operating procedures) across formations in near-real time.
- Red force reasoning: Systems trained to understand and anticipate adversary doctrine and deception techniques.

The result is a force that does not merely react faster—it learns faster, adapts smarter, and maintains initiative even in degraded conditions.

NATO Interoperability and Sovereign Flexibility

Scaling also requires cross-force and multinational compatibility. This can only be achieved through:

- Common data and logic standards—ensuring that allied AI systems can cooperate without compromising national control.
- Modular, sovereign deployments—allowing each state to embed its own doctrine while participating in combined operations.

This is how the UK, France, Germany, Poland, and the Nordics can create an AI-enabled coalition advantage, without surrendering strategic autonomy to foreign platforms or commercial monopolies.

Closing the Kill Chain with Human Dignity Intact – Rethinking Lethality in the Age of Algorithms

Military history has always been shaped by technological overmatch—from the longbow to the rifle, from radar to precision-guided munitions. Yet no advance has so fundamentally challenged the human role in lethality as the integration of artificial intelligence. In this chapter, we address a critical tension: how to accelerate the kill chain without losing sight of what makes war a uniquely human endeavour.

The Kill Chain, Reimagined

The modern kill chain—Find, Fix, Track, Target, Engage, Assess—can be executed in minutes or seconds when powered by AI. Each phase benefits from automation:

- Find & Fix: Sensor fusion, anomaly detection, ISR triage
- Track & Target: Pattern recognition, predictive modelling, real-time prioritisation
- Engage & Assess: Weapon selection, strike coordination, battle damage estimation

Yet, compressing time must not compress accountability. The ethical centre of gravity must remain with the human. A well-functioning AI system should not be measured solely by its kill count—but by its ability to elevate human judgement, reduce collateral harm, and preserve the moral clarity of lawful combat.

> "Machines don't understand the Geneva Conventions. People do." — Lt. Gen. Jack Shanahan, former Director, DoD Joint AI Center

Find & Fix	Track & Target	Human Decision Point	Engage & Assess
AI-assisted sensor fusion with human oversight	Pattern recognition with ethical constraints	Critical human judgment and authorization	Execution with continuous human monitoring

The Lethal Creep of Autonomy

The line between decision-support and autonomous action can blur rapidly. For instance:

- A targeting AI may recommend a strike in a high-threat environment.
- Operators, under time pressure, may increasingly rely on that recommendation without review.
- Over time, the "recommendation" becomes de facto delegation.

Without robust design and doctrinal safeguards, the machine's output becomes the decision, eroding the core principle of human responsibility.

Embedding Human Judgement in the Loop

There are clear, actionable ways to mitigate this risk:

1. Human–on–the–Loop Design

Systems must be structured so that humans can intervene or override, with minimal friction. AI must be interruptible by design—not just in theory.

1. Ethical Failsafes

Pre-set engagement constraints based on IHL (International Humanitarian Law), ROE (Rules of Engagement), and mission-specific doctrine should be encoded at the logic layer, not left to downstream interpretation.

1. Cognitive Guardrails

Alerting systems that flag anomalies, potential errors, or "overconfidence" conditions in AI recommendations should be standard. These act as digital instincts, prompting human scrutiny when most needed.

1. Post-Action Audit Trails

Every lethal action involving AI support must be traceable. Not only for legal accountability but for operational learning and trust-building.

Avoiding Algorithmic Atrocity

History warns us of what happens when lethal technologies outpace human oversight—from the use of chemical weapons in WWI to the automated defences that have misfired in modern conflicts. In the AI era, the greatest threat is not evil intent, but fragile systems making confident mistakes.

Therefore, the emphasis must shift:

- From autonomous kill chains,
- To AI-enabled kill chains with deliberate human custody of force.

Lethality will remain a necessary feature of military force. But in the age of AI, restraint, not speed, may become the highest expression of strength.

Conclusion & Call to Action

As the global strategic balance shifts and the West faces a future of contested domains, degraded alliances, and peer adversaries, it is clear that decision-making—not just firepower—will define military dominance. Human–AI collaboration is not a luxury. It is the precondition for surviving and winning in 21st-century warfare.

Yet too much of the current debate remains trapped between utopianism and paralysis—between the fantasy of full autonomy and the fear of ethical misstep. This white paper has argued for a different path: AI not as a replacement for the human warfighter, but as a trusted cognitive partner embedded at the tactical edge.

> "The real opportunity is not smarter machines—but smarter humans, augmented by machines."

— Paul Daugherty, Human + Machine

Done right, this collaboration will:

- Accelerate decision-making under extreme pressure while keeping humans in command.
- Embed logic, learning, and adaptability into the fabric of every tactical and strategic action.
- Preserve sovereign control and moral responsibility, even as AI speeds up the tempo of combat.

But this vision demands urgent action. Not just in technology development—but in doctrine, procurement, training, and alliance planning.

Call to Action for Stakeholders

For Defence Ministries and Agencies

- Accelerate the development of field-trialled, human-centred AI systems deployable in degraded environments.
- Prioritise funding for cognitive doctrine reform and the operational frameworks required to integrate AI into real-world missions.

For Industry and Contractors

- Move beyond lab demonstrations to rugged, mission-specific AI platforms that empower soldiers, not replace them.
- Build systems that are explainable, interruptible, and sovereign by design, aligned with the operational realities of warfighters.

For NATO and Allied Structures

- Establish a Coalition Human–AI Interoperability Initiative, focused on shared logic layers and trusted teaming protocols.
- Embed human-machine teaming principles into joint training, wargaming, and doctrine development.

The window for shaping the future of human–AI combat collaboration is closing fast. If the West does not define and deploy these systems on our terms, others will do so on theirs—with different values, rules, and strategic aims.

Let us ensure the next evolution in warfare is guided by judgement, fuelled by intelligence, and grounded in the dignity of command.