

Agentic Warfare—Evolving Doctrine

The Problem: Agentic artificial intelligence (AI) and autonomous systems do not just improve existing processes. They change the terms of competition by altering how the force should sense, decide, coordinate, and act. To exploit them, the Department of War (DOW) must rethink concepts, doctrine, command relationships, staff structures, training, force design, acquisition, and the balance between human judgment and machine speed at every echelon. Today, those pieces are moving, but not yet together. Issued in August 2023, the Joint Warfighting Concept 3.0 begins to position DOW for advanced AI and autonomy concepts of operation. Even more authoritative joint warfighting concepts and doctrine are needed. Without them, we are at risk of layering extraordinary new technologies onto an outdated way of war: faster tools, but not faster adaptation; better models, but not better command; more experimentation, but too little change in how the force actually fights.

Why It Matters Now: AI is fast becoming the operating logic of modern warfare. Yet technologies only become decisive when they are integrated with existing combat power through new doctrine, concepts, experimentation, and culture. Tanks mattered when they were fused into combined arms. AirLand Battle mattered because it reorganized how the force thought and fought, not just what it bought. The same is now true of agentic and autonomous systems. Advantage will not go to the military with the most impressive models, but to the one that reorganizes around human-machine teaming fastest and turns machine speed into decision advantage. Our adversaries, especially the People's Liberation Army, are already moving in that direction, integrating AI, autonomy, and data into command, sensing, targeting, and battle networks. If we fail to make the conceptual and doctrinal leap now, we aren't just fielding technology too slowly—we risk entering a machine-speed fight with a dial-up mindset, and in so doing cede decision advantage to our adversaries.

Next Steps: Serious lines of effort are underway to build the foundations of AI-enabled warfare, including the Combined Joint All-Domain Command and Control (CJADC2) system, the 2026 AI Strategy's Pace-Setting Projects, service-specific concept work, and joint- and command-level experimentation. The priority now is convergence: defining a new logic of command, codifying how humans and machines fight together, and using existing concepts, doctrine, and experimentation mechanisms to drive change across the force.

1 - Publish an Authoritative Joint AI Warfighting Concept while also Developing from the Bottom-Up: Building on the Joint Warfighting Concept 3.0, CJADC2, and ongoing work within the Services, this capstone concept should define how commanders, staffs, agents, autonomous systems, data, battle networks, and legacy platforms sense, decide, coordinate, and act together in contested conditions. It should frame AI-era command as a socio-technical warfighting system and anchor a doctrine stack for human-machine command under uncertainty, deception, and degraded networks. The Vice Chairman of the Joint Chiefs of Staff (VCJCS) should sponsor it; the Joint Staff J-7 Directorate should lead it; and the Chief Digital

and Artificial Intelligence Office (CDAO) and the J-6 Directorate should set common AI, data, and network assumptions. Finally, the Services should translate it into doctrine, training, and force design. At the same time, warfighting concept experimentation should continue at lower levels in all key kill chains that will benefit from AI tools, especially cueing, discriminating, targeting, and course of action (COA) generation. This approach will ensure knowledge generated in specific user communities and top-down guidance can co-evolve.

2 - Codify Graduated Human-Machine Authority in Doctrine: AI will increasingly shape what commanders see, what options they consider feasible, and how quickly they can impose tempo. Doctrine must move beyond generic “human in/on the loop” formulas to define different authority models (advise, recommend, orchestrate) graduated by function, echelon, and consequence. Planning, intelligence, surveillance, and reconnaissance fusion, logistics, and battle management can absorb more machine support sooner. Lethal, kinetic, politically sensitive, or strategically escalatory functions require a higher burden of human judgment and legal assurance. Doctrine should also distinguish between connected headquarters, where richer agentic workflows are realistic now, and denied and degraded environments, where narrower, edge-computing tools and clear fallback modes matter more. While Joint Staff J-7 should lead this effort in conjunction with CDAO, in practice, the experiences with AI agents at lower execution levels will provide the key learnings for doctrine.

3 - Turn Experimentation into a Living-Loop AI Warfare Doctrine Production Engine: AI adaptation will fail at force level if lessons stay trapped in pilots, slide decks, and after-action reports. Doctrine must evolve as quickly as operational practice. The missing piece is a fixed pipeline that turns insights from AI-enabled exercises and experiments into not only improvements of the AI tools themselves, but also rapid doctrinal outputs. Exercises and experiments should result in Joint Doctrine Notes, digital playbooks, authority matrices, and updates to tactics, techniques, and procedures related to AI, while building on existing mechanisms like joint doctrine processes, the Air Force's Decision Advantage Sprint for Human-Machine Teaming (DASH), and Project Convergence. Those events must stress AI under real adversary conditions, from jamming and spoofing to cyber attack, data poisoning, and deception. That would connect experimentation to force development, training, and program change. The VCJCS should sponsor it; Joint Staff J-7 should lead it; and the combatant commands, Service battle labs, CDAO, and J-6 should feed it.

4 - Rebuild Command and Control for the Age of AI: AI adoption will stall, or even become dangerous, if commanders and staffs do not know how to direct these systems, challenge their outputs, recognize their failure modes, and keep fighting when they degrade. Commanders must understand how to lead hybrid human-machine teams without surrendering judgment or becoming dependent on algorithms. Exercises and experimentation events should enable commanders to try different levels of “risk taking” to find the balance between trusting AI agents and using human judgment. Operator trust, shared understanding, and command performance under degraded conditions should be treated as measures of combat readiness.