READY, FIRE, AIM

Tactical Autonomy in the Age of AI

S. Tucker Browne
D. Isaiah Harp
Michael W. Byrnes
Brent A. Maggard

This article critically examines the Air Force's strategic pivot toward a future force heavily reliant on tactical autonomy. Drawing lessons from other technical fields, this article identifies three fundamental problems—perception, data, and adversarial vulnerability—that undermine the feasibility of autonomous combat aircraft and threaten the Air Force's allocation of resources, operational effectiveness, and long-term advantages in air superiority. To achieve lasting strategic gains, the service must immediately reassess future investments and planning with rigorous technological realism, focusing on verifiable performance, validated operational concepts, and resilience against adversarial countermoves. The realignment of future force planning with technological reality can be accomplished by measures focused on realistic capability demonstration, disciplined procurement, and strategic hedging.

In June 2023, US Army General Mark Milley, then-chairman of the Joint Chiefs of Staff, affirmed what a rising crescendo of public and private actors had already observed: the world was witnessing "the most fundamental change" in the history of the character of war, including "the introduction of robots," "a pilotless Air Force," and artificial intelligence (AI).[1] As with previous paradigm shifts from muskets to rifles or conventional to nuclear weapons, AI would forever divide military history into a distinct before and after.

Major S. Tucker Browne, USAF, serves as a program element monitor with the Office of the Assistant Secretary of the Air Force for Acquisition, Technology, and Logistics. He holds a master of science and technology intelligence from the National Intelligence University, a master of science in flight test engineering from the US Air Force Test Pilot School, and a master of engineering from Cornell University.

Major D. Isaiah Harp, USAF, is a fellow at the Institute for Future Conflict and an instructor in the department of aeronautics at the US Air Force Academy. He holds a master of science in aeronautical and astronautical engineering from Stanford University and a master of science in flight test engineering from the US Air Force Test Pilot School.

Lieutenant Colonel Michael Byrnes, USAF, PhD, is the commander of the 732nd Operations Support Squadron, Creech Air Force Base, Nevada.

Major Brent Maggard, USAF, the command chief pilot for the F-22 developmental test, Air Dominance Combined Test Force, holds a master of science in flight test engineering from the US Air Force Test Pilot School.

1. Jim Garamone, "Milley Makes Case for Rules-Based Order, Deterrence in New Era," US Department of Defense (DOD), 30 June 2023, https://www.defense.gov/.

Such a shift is reflected in the Pentagon's strategic plans. The Air Force's vision for air superiority in the 2030s provides a particularly clear case study of how a major acquisition program has evolved adjacent to the apparent revolutions in AI. In the 2010s, the Next-Generation Air Dominance program emerged with a plan to deliver a sixth-generation manned fighter in the 2020s.[2] By 2022, it had developed into a "family of systems," emphasizing "much less expensive autonomous uncrewed combat aircraft" known as collaborative combat aircraft (CCA).[3] In July 2024, the service paused the program's manned fighter component, partially because of the platform's cost but also ostensibly due to "technology advances" in autonomous systems.[4] But after months of debate, in March 2025, the program was once again reinstated.[5]

The Air Force remains entrenched in an internal struggle to determine the extent of the role that autonomy should play in its future force structure. While some senior leaders and technologists question the feasibility and strategic wisdom of fully autonomous tactical aircraft, a prominent contingent continues to advocate for rapid advancement toward a largely autonomous combat force.[6] A review of Air Force budget requests for autonomous combat aircraft programs shows how planned funding for autonomous aircraft development dramatically increases in 2026 to 2029 (fig. 1).[7]

The logic behind this strategic pivot is compelling. If an autonomous aircraft can perform the tasks of a human fighter pilot at the leading edge of combat—providing tactical autonomy—then AI-powered wingmen can mitigate the numerous disadvantages of human pilots. Relatively inexpensive aircraft without human operators can absorb tactical risk through attrition or distraction. Each robot will be as skilled as all other robots, and its software can infinitely reproduce new skills. The lengthy and costly enterprise of training human capital will be reduced to a copy-and-paste operation. Advanced AI systems may even generate novel solutions to tactical problems that humans have never imagined.[8] If the Air Force is on a credible path to tactical autonomy, then it is imperative to proceed with total commitment toward this potential offset.

2. Aaron Mehta, "Kendall Unveils 6th Gen Fighter Strategy," *Defense News*, 1 February 2015, https://www.defensenews.com/.

3. Charles Pope, "Kendall Details 'Seven Operational Imperatives' & How They Forge the Future Force," US Air Force [USAF, website], 3 March 2022, https://www.af.mil/.

4. John Tirpak, "CCA Contract Expected in Fall; First Versions Under Construction," *Air & Space Forces Magazine*, 6 July 2024, https://www.airandspaceforces.com/.

5. Matthew Olay, "Trump, Hegseth Announce Air Force's Next Generation Platform," DOD, 21 March 2025, https://www.defense.gov/.

6. Audrey Decker, "Robot Reality Check: Crewed Warplanes Will Remain Vital for Years, USAF General Says," *Defense One*, 7 December 2024, https://www.defenseone.com/.

7. USAF Financial Management and Comptroller, *Department of Defense Fiscal Year (FY) 2023 Budget Estimates: Air Force*, vol. 2, Research, Development, Test & Evaluation (Department of the Air Force [DAF], April 2022), https://www.saffm.hq.af.mil/; and *Department of Defense Fiscal Year (FY) 2025 Budget Estimates: Air Force*, vol. 2, Research, Development, Test & Evaluation, Air Force (DAF, March 2024), https://www.saffm.hq.af.mil/.

8. Daniel Castro and Joshua New, *The Promise of Artificial Intelligence* (Center for Data Innovation, October 2016), https://www2.datainnovation.org/.

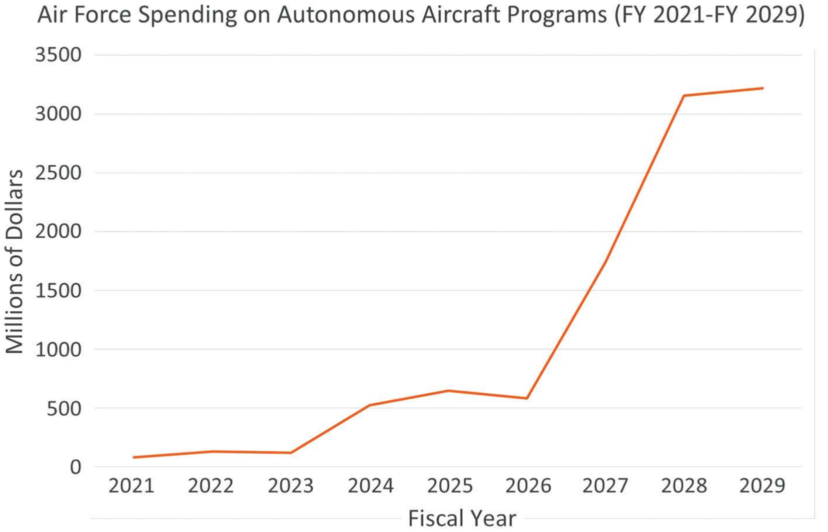

Figure 1. Annual Air Force spending on autonomous aircraft programs, fiscal years 2021 to 2029, past and projected

Unfortunately, however, the Air Force is not on such a path. The service has not acknowledged, accounted for, or mitigated three fundamental issues that render this vision not just impractical but also fantastical in the near-term: the problems of perception, data, and adversarial vulnerability.

First, even advanced autonomous systems routinely fail to accurately perceive well-characterized, relatively benign environments. Combat environments are inherently complex and uncertain, and AI systems have yet to demonstrate the sense-making required to build reliable situational awareness. Second, autonomous systems struggle to generalize beyond their training data set. Because war is the province of uncertainty and intelligence of the adversary is always incomplete, the military will not be able to produce the training data required for air combat in the pre-hostilities time frame—precisely when needed. Finally, AI systems exhibit precarious brittleness when faced with even rudimentary adversarial attacks.

These persistent, interrelated barriers threaten the Air Force's allocation of resources, operational effectiveness, and long-term advantages in air superiority. The service must immediately reassess its planned investments with rigorous technological realism, focusing on verifiable performance, validated operational concepts, and resilience against adversarial countermoves. Such a disciplined approach will ensure durable strategic gains while avoiding hollow capabilities that could emerge from a force design reliant on unproven autonomous systems.

Ready, Fire, Aim

Combined, the three issues of perception, data, and adversarial vulnerability suggest that the Air Force's pivot to tactical autonomy is a case of premature action akin to "Ready, Fire, Aim." These challenges are not mere technical hurdles that time, effort, and funding will inevitably overcome but core limitations of a family of technologies whose theoretical foundations remain incomplete.[9] If the Air Force pulls the trigger now, it risks missing wildly. Furthermore, considerations of a staged approach to adoption—such as the Air Force's recent shift to tethering autonomous aircraft to nearby piloted platforms—fundamentally misunderstand the nature of these challenges.

This analysis does not seek to dismiss the potential of AI in military aviation. As test pilots with thousands of collective hours flying remotely piloted aircraft and stealth fighters in simulated and actual air combat, the authors fully believe in technologically advanced warfare in which computational systems and the human mind combine to violent advantage. Yet, a warfighting organization must not stake its survival on an optimistic future where ever-increasing data and compute enable "models [to] simply get better" forever.[10]

Additionally, the authors do not advocate for preservation of the status quo. While the Air Force's current force structure has proven capabilities, it must evolve to meet the challenges posed by emerging threats. The primary concern is that an aggressive pivot toward fully autonomous tactical aircraft represents a premature and high-risk path compared to other solutions—one that overlooks critical technology limitations that render tactical autonomy unrealistic in the near-term. In an era of strategic competition, such undisciplined technological optimism at institutional scale presents national risk.

Lessons from Other Fields

A critical examination of how similar optimism regarding AI has played out in other technical fields clarifies the scope of these risks. A recurring pattern of AI booms and winters has repeated continuously since the 1950s, and this should serve as a warning for senior Air Force leaders.[11] First, early AI successes in controlled environments generate widespread enthusiasm. Next, enthusiasm fuels bold predictions and significant investment. When these systems are tested in the messy reality of the real world, they encounter unexpected limitations or technical obstacles that prove insurmountable. Finally, performance stalls well below the level needed to realize revolutionary change.

IBM's computer system Watson exemplified this cycle. In 2011, Watson defeated human champions on the gameshow *Jeopardy!*, leading to predictions that Watson would revolutionize complex decision-making across industries, particularly in healthcare. But as the program manager stated, "The challenges turned out to be far

9. J. Mark Bishop, "Artificial Intelligence Is Stupid and Causal Reasoning Will Not Fix It," *Frontiers in Psychology* 11 (2021), https://doi.org/.

10. Charlie Warzel, "AI Has Become a Technology of Faith," *The Atlantic*, 12 July 2024, https://www.theatlantic.com/.

11. Amirhosein Toosi et al., "A Brief History of AI: How to Prevent Another Winter (a Critical Review)," Cornell University, arXiv, 1 October 2021, https://ar5iv.labs.arxiv.org/.

more difficult and time-consuming than anticipated."[12] In 2022, after investing over $4 billion, IBM sold Watson Health for a fraction of its cost.[13] While it excelled at highly structured knowledge retrieval tasks, Watson ultimately could not achieve its goal of meaningfully contributing to medical diagnoses in the real world.

The autonomous vehicle industry offers even closer parallels to Air Force ambitions. Since 2010, firms have spent over $100 billion on autonomous vehicle technology development.[14] This figure dwarfs all current and projected military spending on AI, yet the industry produced surprisingly modest, incremental successes contrasted to intended outcomes.[15] Companies like Waymo and Cruise, despite operating only in carefully mapped cities, have yet to successfully deploy fully autonomous vehicles and endure multibillion dollar deficits annually.[16] As of 2024, Tesla—despite its bold claims since 2016 that "Full Self-Driving" was "just around the corner"—remains only at Level 2 of 5 autonomy, requiring constant driver supervision and intervention.[17] Given a relatively structured problem set, training data measured in millions of hours, and predictable traffic rules, these products still encounter crippling difficulties navigating anomalous events, bad weather, unpredictable humans, or left-hand turns.[18]

AI proponents argue that solutions are forthcoming, pointing to ongoing methodological improvements and increased data collection. Yet the struggle revolves around a fundamental challenge of AI: it can excel in structured, predictable environments but degrades rapidly in "edge cases"—situations where the underlying problem diverges from the expected model or training data coverage. As cognitive psychologist Steven Pinker observed, the main lesson of AI research is that "the hard problems are easy, and the easy problems are hard."[19] If AI-based vehicles struggle to safely operate autonomously on well-mapped city streets—a task entrusted to some 16-year-olds—how can the military expect AI to perform within the chaotic and hostile environment of real-world aerial combat? There, decisions are replete with uncertainty, training data are scarce, and the adversary consents to few rules.

---

12. Steve Lohr, "What Ever Happened to IBM's Watson," *The New York Times*, 18 July 2021, https://www.nytimes.com/.

13. Clare Duffy, "IBM Is Selling Off Its Watson Health Assets," CNN, 21 January 2022, https://www.cnn.com/.

14. Max Chafkin, "Even After $100 Billion, Self-Driving Cars Are Going Nowhere," *Bloomberg*, 6 October 2022, https://www.bloomberg.com/.

15. B. Padmaja et al., "Exploration of Issues, Challenges and Latest Developments in Autonomous Cars," *Journal of Big Data* 10, no. 1 (2023), https://doi.org/.

16. Trisha Thadani, "Embattled Self-Driving Car Company Cruise Lost $3.48 Billion in 2023," *The Washington Post*, 31 January 2024, https://www.washingtonpost.com/.

17. Ty Duffy, "What Tesla Autopilot and Full-Self Driving Can and Can't Do," *InsideEVs*, 27 November 2024, https://insideevs.com/.

18. Chafkin, "Self-Driving Cars."

19. Steven Pinker, *The Language Instinct* (DA Information Services, 1994).

Operational analyses frequently assert that autonomous combat systems will bring value by simply increasing the number of sensors and shooters in the battlespace.[20] Using that logic, the onboard autonomy can be initially rudimentary; the mere presence of additional "iron in the sky" enhances overall combat capability. Yet this perspective overlooks a critical lesson from other fields: quantity does not compensate for quality in perception and sense-making. Adding more sensors does not grant each robot a more accurate understanding of the battlespace, any more than adding more cameras makes a better self-driving car. Additional shooters provide little advantage if they cannot both construct and build situational awareness of the battlespace composition and context.

An AI optimist might sidestep this reality and instead imagine that autonomous combat aircraft will receive a clear, reasonably reliable picture of the battlespace. The AI-managed task would then be to optimize tactical decisions based on this high-quality information. This view reflects a misunderstanding of air combat. A common operating picture that fuses trustworthy information and presents it reliably does not exist. Rather, human operators must construct their understanding of reality from a tangled web of conflicting and incomplete data sources. Sensor displays may present misleading information, sometimes showing friendly aircraft as enemies, enemy aircraft as friendlies, phantom tracks where no aircraft exist, and empty spaces where aircraft actually fly. Adversaries exacerbate these challenges as they actively seek to deny, degrade, disrupt, and deceive every aspect of battlespace awareness.

Experienced pilots navigate perceptual uncertainty through an intricate combination of inference, reasoning, contextual awareness, and tactical intuition honed over years of training. They can recognize when information is unreliable, adjust their mental models accordingly, and adapt their tactics in real time. Perhaps most importantly, when a pilot knows they do not understand the situation, they can acknowledge this uncertainty, seek additional information, and take novel actions to mitigate risk—distinctly human behaviors that lead to reasonable decisions in the face of incomplete or misleading information.

Viewed naively, AI successes in games such as chess and Go seem to evidence future capability. These successes, however, are attributable to the fact that AI bypasses rather than engages with these perception challenges. Most game-playing AI agents are given perfect or near-perfect information about the game. Additionally, the space of all possible decisions for these agents is both discrete (players must move on the spaces) and finite (players must follow the rules). In combat, the decision space is continuous and infinite.

Understanding the context of previous AI successes and failures is crucial for understanding why the three critical problems of perception, data, and adversarial vulnerability are closer to insurmountable barriers than mere technical hurdles. The

---

20. "Collaborative Combat Aircraft for Disruptive Operations Mitchell Institute CCA Wargame Executive Summary," slides, Mitchell Institute for Aerospace Studies, 2024, https://www.mitchellaerospacepower.org/.

pattern observed in other fields, where AI systems hit hard limitations that resist incremental solutions, is likely to become even more pronounced for air combat. Furthermore, their combined effects create challenges that surpass anything encountered in medical diagnosis, autonomous vehicles, or complex games.

Three Critical Problems

Understanding the three challenges to tactical autonomy and how they mutually amplify each other indicates how current AI technologies will fail to deliver the autonomous combat aircraft hyped by proponents.

Perception in Complex Combat Environments

Even state-of-the-art autonomous systems routinely fail to accurately perceive well-characterized, familiar human environments.[21] A 2023 study found that ineffective perception and sensing in off-nominal environmental conditions has "been the problem that keeps autonomous vehicles from going to higher autonomy." The study faulted every aspect of autonomous perception systems: the sensors, the fusion algorithms, and the AI.[22] In 2019, the National Transportation Safety Board determined that a Tesla driver was decapitated in a crash because his "Autopilot vision system did not consistently detect and track" a broadside tractor-trailer as an object or threat.[23] The report showed that the tractor-trailer was continuously visible to the human eye five seconds prior to the collision; however, all the way through the moment of impact, the car never braked nor steered.

Combat environments are significantly more complex and require sensors to tackle the challenge of building situational awareness across hundreds of miles. Yet recent autonomy programs such as the Defense Advanced Research Projects Agency (DARPA) AlphaDogfight Trials have elided the problem of perception entirely. In these trials, an AI agent defeated an experienced F-16 pilot in a one-on-one simulation of a task known as basic fighter maneuvers.[24] But, critically, the AI software had perfect information about the simulated environment, capabilities of both aircraft, and the adversary's position, speed, and direction in real time; the human had no such advantages.[25]

---

21. See Ashley Roque, "Frustrations Mount over Army's Robotic Combat Vehicle Autonomy, Acquisition Approach," *Breaking Defense*, 22 July 2024, http://breakingdefense.com/; and Mary L. Cummings, "What Self-Driving Cars Tell Us About AI Risks," *IEEE Spectrum*, 30 July 2023, https://spectrum.ieee.org/.

22. Yuxiao Zhang et al., "Perception and Sensing for Autonomous Vehicles Under Adverse Weather Conditions: A Survey," *ISPRS Journal of Photogrammetry and Remote Sensing: Official Publication of the International Society for Photogrammetry and Remote Sensing (ISPRS)* 196 (2023), https://doi.org/.

23. *Highway Accident Brief: Collision Between Car Operating with Partial Driving Automation and Truck-Tractor Semitrailer* (National Transportation Safety Board, 22 January 2020), https://www.ntsb.gov/.

24. "AI Bests Human Fighter Pilot in AlphaDogfight Trial at Johns Hopkins APL," Johns Hopkins University Applied Physics Laboratory, press release, 28 August 2020, https://www.jhuapl.edu/.

25. Adrian P. Pope et al., "Hierarchical Reinforcement Learning for Air-to-Air Combat," arXiv, last revised 11 June 2021, http://arxiv.org/.

In another example, after then-Secretary of the Air Force Frank Kendall flew in the X-62—an experimental AI test-bed performing basic fighter maneuvers as part of DARPA's follow-on Air Combat Evolution program—he stated that the AI had "roughly an even fight" with an experienced fighter pilot.[26] The Air Force press release declared that "the controls of the X-62A remained untouched by both Kendall and the safety pilot in the backseat throughout the entire test flight."[27] But this claim was not true; the AI software was only activated in specific bounded portions of the test flights. In reality, the X-62 could not autonomously take off, transit to the airspace, set up for the dogfight, dogfight without the adversary aircraft sharing an uninterrupted feed of high-quality information, safely end the dogfight, safely return to base, nor land on its own. When researching these claims, the authors contacted Air Force public affairs to ask about the article's overstatements. In response, public affairs corrected their press release.[28]

While ongoing experiments like AlphaDogFight and Air Combat Evolution are interesting lines of research, they bypass real machine perception challenges critical to tasks such as basic fighter maneuvers and focus solely on AI maneuver geometry problem-solving—more akin to playing a videogame than evaluating realistic tactical problems. In combat, adversaries will not provide high-quality feeds of their positions. Air Force discussions about the maturity of autonomy should not equate narrow unrealistic applications of AI that entirely bypass sense-making to broad tasks fighter pilots perform in combat.

While advances in computer vision algorithms can seem impressive, still-struggling outcomes in autonomous vehicle applications, coupled with an absence of research on how these classifiers perform on military sensors, suggest a long road ahead for even basic object classification in air combat. Assuming object classification challenges were solvable, a massive research gap would still remain between detecting or labeling objects and sense-making of the resulting, inevitably imperfect, tactical picture.

Limited Training Data for Real-World Conflict

The second problem for tactical autonomy is the scarcity of useful datasets to train the systems. While autonomous systems trained on plentiful, accurate, and well-labeled data show promising results when applied to narrow problems, they struggle or fail to generalize results outside their training data set. A complete training data set for air combat is impossible to produce in peacetime.

---

26. Jon Harper, "Air Force's Kendall: AI Agents Had 'Roughly an Even Fight' Against Human F-16 Pilot in Recent Engagements," *DefenseScoop*, 8 May 2024, https://defensescoop.com/.

27. "27 Nov 24 Archive: Air Force's Kendall: AI Agents Had 'Roughly an Even Fight' Against Human F-16 Pilot in Recent Engagements," Wayback Machine: Internet Archive, 3 May 2024, accessed 10 December 2024, https://web.archive.org/.

28. Gary Hatch and Mary Kozaitis, "SecAF Kendall Experiences VISTA of Future Flight Test at Edwards AFB," USAF, 3 May 2024, https://www.af.mil/.

The last two documented US air-to-air kills against manned fighters occurred in 1999 and 2017, highlighting the scarcity of air combat data in the twenty-first century.[29] Moreover, much of the available military data on air combat tactics and outcomes is classified, fragmented, unlabeled, or not representative of current technology. A 2019 RAND Corporation report highlighted that significant portions of available military data are not stored in accessible formats, lack interoperability, and are often not understandable or traceable, all of which exacerbate these problems.[30]

Proponents of autonomous combat systems often suggest that synthetic (computer-generated) data and advanced simulations can overcome the training data shortage. This argument misconstrues both the nature of modern air combat and the limitations of simulated environments. Modern air combat involves complex interactions between physical materials, electromagnetic waves, the correlation of multiple sensor feeds, and datalinks—physics interactions that are exceedingly difficult to model accurately. A United Nations policy paper highlights the risk in this approach: "Poorly generated synthetic data can lead to inaccurate and unreliable AI models."[31] More fundamentally, synthetic data only incorporates known variables and interactions. Real combat presents scenarios never considered in training or predicted via intelligence collection. Unlike games with fixed rules, combat tactics continuously evolve as adversaries create scenarios outside expected parameters. While automated domain randomization can increase the size of training datasets, it still fails to account for edge cases, adversarial ingenuity, and intelligence uncertainties.

The relevance of existing datasets also diminishes with the introduction of new warfare technologies. An AI trained on data from five years ago would find itself wholly unprepared for the realities of combat today. While a human can easily take training from old technology and tactics and update it with new assessments, the failure of contemporary AI to adapt to even trivial "distribution shifts" of the training data has been demonstrated repeatedly. A 2021 study tested the transition of an automated breast cancer detection algorithm from one hospital to another, while keeping every anticipated factor unchanged. Its performance dropped from 93 to 70 percent based on unexpected confounding factors such as the new hospital's lighting, patient demographics, and photography procedures.[32]

29. Oriana Pawlyk, "US F/A-18E Shoots Down Syrian Su-22 in Air-to-Air Kill," Military.com, 18 June 2017, https://www.military.com/.

30. Danielle C. Tarraf et al., *The Department of Defense Posture for Artificial Intelligence* (RAND Corporation, 17 December 2019), https://www.rand.org/.

31. Philippe de Wilde et al., *Recommendations on the Use of Synthetic Data to Train AI Models* (UN University, 29 February 2024), https://unu.edu/.

32. Pang Wei Koh et al., "WILDS: A Benchmark of in-the-Wild Distribution Shifts," paper presented at the 38th International Conference on Machine Learning, Vienna, Austria, 13–18 July 2020, https://cs.stanford.edu/.

Similarly, highly structured air combat training events such as Red Flag will not provide the density or scale of training data required to realize tactical autonomy.[33] In training, instructors of wingmen do not aim to show them everything they could possibly see in combat. They train them on how to deal with the previously unseen and the unexpected. A human can be taught to react and improvise, whereas the most advanced algorithms today can only regurgitate training data in complex ways. As one expert observes, in seven decades AI researchers have made "almost no progress" in "apply[ing] knowledge from one domain to another."[34]

Overall, accurate and plentiful data representative of realistic combat scenarios does not exist and cannot be easily simulated. The challenge is not one of quantity or labeling. AI systems that cannot generalize beyond their flawed training will fail in combat. They will overfit or underfit to limited, imperfect data and make critical errors when faced with novel situations. More critically, this data limitation ensures that any perceptual capabilities the AI develops will be inherently flawed, creating vulnerabilities that adversaries can exploit.

Fragility to Adversarial Attacks

Today's AI systems are unacceptably brittle—vulnerable to catastrophic failure—when faced with adversarial attacks.[35] In 2017, scientists showed AI's fragility in the realm of image recognition.[36] In an attack on autonomous driving systems, researchers applied small stickers to road signs that caused an image recognition algorithm to identify a stop sign as a speed limit sign with over 90 percent confidence. In 2024, studies revealed that attacks like these could also be highly effective in the military domain. Using "black-box" techniques—attacks that were not reliant on exploitation or understanding of the algorithms they were attacking—researchers demonstrated that imperceptible "universal adversarial perturbations" could cause AI systems designed to recognize targets in synthetic aperture radar imagery to misclassify military vehicles such as tanks up to 64 percent of the time.[37]

Alarmingly, these AI systems under attack often report high confidence in their incorrect decisions, providing no indication that they have been compromised.[38] This overconfidence stems directly from the first two problems: systems with flawed per-

33. "414th Combat Training Squadron 'Red Flag,' " Nellis Air Force Base, current as of October 2022, https://www.nellis.af.mil/.

34. Bishop, "Artificial Intelligence."

35. Katherine Tangalakis-Lippert, "Marines Fooled a DARPA Robot by Hiding in a Cardboard Box While Giggling and Pretending to Be Trees," *Business Insider*, 29 January 2023, https://www.businessinsider.com/.

36. Kevin Eykholt et al., "Robust Physical-World Attacks on Deep Learning Models," arXiv, last updated 10 April 2018, http://arxiv.org/.

37. Bowen Peng et al., "An Empirical Study of Fully Black-Box and Universal Adversarial Attack for SAR Target Recognition," *Remote Sensing* 14, no. 16 (2022), https://doi.org/.

38. Jingshu Li and Yitian Yang, "Overconfident and Unconfident AI Hinder Human-AI Collaboration," arXiv, 12 February 2024), https://arxiv.org/.

ception, trained on limited data, cannot recognize deception, creating a critical vulnerability in combat environments where rapid, accurate decision-making is essential.

Recent examples demonstrate the severity of this vulnerability. In 2023 the program KataGo, a Go-playing agent considered "superhuman," was defeated in more than 97 percent of its matches by amateur players who employed adversarial strategies that exploited its inflexibility.[39] That same year, Marines defeated an advanced DARPA AI surveillance system by walking around in a cardboard box, holding branches and "pretending to be trees," or doing somersaults.[40] Discussion around autonomy in the Air Force rarely includes the fact that crude adversarial attacks can often defeat modern AI systems in unexpected or unpredictable ways.

The three problems identified above have created an intractable situation: systems that cannot reliably perceive their environment, trained on insufficient data, become highly vulnerable to relatively trivial adversarial manipulation. Entrusting autonomous systems with control in combat scenarios embeds unmitigated systemic risk within those operations. This risk, which stems from the immaturity of the technology and its incomplete theoretical grounding, has not been adequately considered or accounted for in either force planning assumptions or operational analyses used to make major decisions. Given these formidable technical challenges, it is worth examining how current mitigation approaches address—or fail to address—these fundamental barriers.

Tethering: An Incomplete Solution

The Air Force's position on the degree of tactical autonomy required to achieve operational advantage has continuously shifted. In early 2023, senior leaders emphasized that CCA must operate "untethered with a high level of autonomy" to function in a "contested electromagnetic spectrum."[41] By late 2024, they reversed course, stipulating that CCAs "have to be under tight control" with "line-of-sight communications." The Secretary of the Air Force stated that "the default, if [CCA] lose communications, would be for them to return to base, which takes them out of the fight."[42] In 18 months, the entire CCA employment concept shifted from high reliance on machine autonomy to complete reliance on tethering, where humans supervise robots from nearby fighter aircraft.

This dramatic shift represents a forced retreat from initial autonomy claims without explicit recognition of their inherent limitations in either strategic communications or acquisition planning. Although proponents may see tethering as a viable long-term

39. Tony T. Wang et al., "Adversarial Policies Beat Superhuman Go AIs," paper presented at the Deep RL Workshop at NeurIPS 2022, 9 December 2022, https://people.eecs.berkeley.edu/.

40. Tangalakis-Lippert, "Marines."

41. Jon Harper, "Air Force Preparing for 'Tethered' and 'Untethered' CCA Drone Operations," *DefenseScoop*, 27 March 2023, https://defensescoop.com/.

42. Michael Marrow, "CCA Drones May Not Be Tied to NGAD, Need Line-of-Sight Control: Kendall," *Breaking Defense*, 16 September 2024, https://breakingdefense.com/.

mitigation, it undermines autonomy's strategic promise by reintroducing human dependency at the point where independence is essential. While tethering might provide a temporary workaround for technology demonstrations and limited experimentation, it is not a replacement for real solutions to the core technical problems facing tactical autonomy systems. As a result, efforts to scale this limited solution into future operational forces fail to resolve the fundamental problems while introducing three new challenges.

First, tethering contradicts the core strategic rationale for autonomous aircraft—the ability to operate independently in contested environments where communication may be severed—while imposing additional tactical workloads on operators. The potential advantages of supervised autonomous aircraft have been explored in simulators, in which computer-driven characters in computer-defined worlds leverage perfect knowledge to behave in advantageous ways.[43] But in the physical world, AI-driven systems will stumble over the three critical problems. The resulting limitations force aircrew to expend valuable mental resources supervising unreliable robotic platforms that range from marginally functional to completely unpredictable. Given that cognitive workloads and stresses in aerial combat are intrinsically extreme, the tradeoff between cognitive costs spent supervising robotic wingmen and their tactical value deserves scrutiny.

Second, whether the institutional Air Force can contend with such rapid shifts in foundational assumptions about autonomous combat aircraft is unclear. How many operational assessments that informed the future force structure assumed the use of untethered autonomous platforms? How many models and simulations that looked at autonomous wingmen locally tethered to fighter aircraft assumed simulated performance where simulated robots did not have to contend with the three critical problems? The conclusions of operational analyses are highly sensitive to the assumptions that drive them.

Finally, tethering creates an obvious vulnerability that competent adversaries will exploit. When a single pilot controls multiple CCAs, destroying or disrupting that pilot's aircraft suddenly removes multiple platforms from the battle. This creates a strong incentive for enemies to focus overwhelming force on the controlling aircraft. The more CCAs each pilot controls, the more attractive the controlling platform becomes as a target. This vulnerability is magnified in contested environments where communications jamming could force CCA formations to automatically retreat or be rendered functionally inoperative. This relationship, the available methods for manipulating it through network design, its impact on enemy tactics, and the investment in communications infrastructure and human-machine interfaces required to mitigate it remain understudied.

While tethering might enable initial CCA experimentation, it cannot serve as the foundation for future autonomous combat systems. An Air Force heavily reliant on

43. Mitchell Institute for Aerospace Studies, "Aerospace Nation: Gen Kenneth S. Wilsbach," 10 July 2024, YouTube video, 59:09, https://youtu.be/.

tethered autonomous platforms would sacrifice the advantages autonomy was supposed to provide while introducing new vulnerabilities. The vision of a future force structure built around autonomous combat aircraft remains fundamentally unrealistic, as the three critical problems represent inherent limitations rather than temporary obstacles awaiting breakthrough solutions.

Risks of Technological Optimism

Analysis of the critical problems facing combat AI reveals a troubling possibility: the optimism surrounding autonomous combat aircraft may be leading the Air Force toward a path that appears promising at first but ultimately leads to a failed end-state requiring a costly reversal. The true cost of such a blunder will be measured not only in the resources directly expended on autonomy development but also in opportunity costs and capability gaps that compound with time.

There is ample evidence of such optimism at work. In December 2023, senior Pentagon leaders stated that in the future CCA could expand to roles such as collaborative reconnaissance or mobility aircraft but that the immediate priority was to "focus on an air-to-air mission."[44] This prioritization inverts the logic of technology development. Semi-autonomous reconnaissance aircraft such as the RQ-4 exist today in large numbers. Autonomous mobility aircraft have recently been demonstrated in Federal Aviation Administration-approved flights.[45] The bulk of tasks these aircraft perform—ground operations, takeoff, cruise, descent, landing—will also have to be performed by CCA. A realistic developmental path would seek to mature existing semi-autonomous aircraft before attempting more complex autonomy tasks. Yet these foundational applications remain only potential concepts, while the Air Force plans for 1,000 CCA as the first step toward fielding tactical autonomy.[46]

The Air Force's plan to leapfrog multiple steps of technology maturation has echoes of the US Navy's littoral combat ship program, which one researcher observes "was essentially counted to solve every single one of the Navy's problems all at once."[47] Yet three years after fielding, then-Chief of Naval Operations Admiral Michael Gilday asked to retire many of the ships because they "did not work out technically." At issue were the technological immaturity of perception subsystems such as radar and sonar, and unreliable engines that could not be repaired underway because the Navy wanted to operate the ships with fewer humans on board.[48] Similar optimism regarding tactical autonomy predominates in the Air Force today, with proponents finding a panacea for a wide range

44. Dave Deptula et al., "Collaborative Combat Aircraft Vectors," transcript, panel discussion, ASC [Air, Space, Cyber] Conference 2023, 11 September 2023, Joint Base Andrews, Maryland, https://www.afa.org/.

45. Mark Phelps, "Successful Remote-Piloted Flight for Cessna Caravan," AVWeb, 6 December 2023, https://www.avweb.com/.

46. Tirpak, "CCA Contract."

47. Oren Liebermann et al., "US Navy Chief Defends Plan to Scrap Troubled Warships Even Though Some Are Less than 3 Years Old," CNN, 12 May 2022, https://www.cnn.com/.

48. Liebermann et al., "Navy Chief."

of problems—from battlespace awareness to the endemic pilot shortage—in a single technologically immature idea.

Heavy investment in immature technology also entails opportunity costs. Every dollar spent in pursuit of tactical autonomy diverts resources from viable, crucial technologies relevant to tactical aviation, including datalinks, long-range weapons, multi-spectral sensors, and piloted platforms. Additionally, the risk of bureaucratic path lock-in is acute, exacerbated by the sunk cost fallacy and institutional inertia. As more resources are invested in tactical autonomy, it becomes both psychologically and politically difficult to change course, even in the face of mounting evidence of the technology's limitations. Large bureaucracies like the military are particularly susceptible to this problem, as careers of both officers and defense contractors become tied to the success of programs that they work on. These forces, if left unchecked, create a self-reinforcing incentive to continue down the combat autonomy path regardless of emerging limitations.

This inertia is exacerbated by a disconnect between public Air Force messaging about autonomous capabilities and the private recognition of limitations among technical experts and program managers. The resulting gap creates an environment where realistic assessments struggle to influence strategic planning.

It is critical not to confuse limited demonstrations of autonomy capabilities with proof of feasibility for broader and more ambitious tactical autonomy goals. Plans for 2030s air superiority that rely on autonomous combat systems stake success on platforms with fundamental operational concepts unproven even in limited testing environments. Given the substantial technical gaps identified, prudent risk management demands the Air Force explicitly define realistic milestones, maintain clear-eyed strategic hedges, and avoid prematurely assuming that incremental successes in canned demonstrations guarantee operationally capable autonomous machines.

The Path Forward

The disconnect between the maturity of technologies underpinning tactical autonomy and the strategic plan for its adoption requires reconciliation. Historically, the service has bought down risk incrementally, balancing the imperative to adapt quickly with the necessity of getting the adaptation right. To realign future force planning with technological reality, the authors advocate for three specific measures focused on realistic capability demonstration, disciplined procurement, and strategic hedging.

First, the Air Force should establish clear, meaningful acquisition milestones that limit tactical autonomy systems from advancing beyond the research, development, testing, and evaluation (RDT&E) phase until they have demonstrated genuine combat utility. These milestones should emphasize warfighter-validated performance against the three critical problems: perception in contested environments, operational effectiveness despite realistic data limitations, and resilience against adversarial tactics. Assessment should blend objective measurements where possible with warfighter evaluation of tactical utility, as these systems must ultimately prove their worth to the operators who will employ them. Tactical autonomy must earn its place in the combat air force via proven performance in realistic combat scenarios, not merely controlled simulations, scripted tests, or projected analyses.

Second, the Air Force should clearly differentiate procurement objectives between early-increment CCAs—semi-autonomous systems requiring human oversight—and future, highly autonomous aircraft intended for operational integration. Procurement of early-increment CCAs should be limited to only the numbers necessary for realistic operational experimentation, technology validation, and tactical concept refinement. An initially smaller fleet of CCAs would enable these tasks without prematurely institutionalizing unproven autonomy assumptions. If developmental autonomy systems demonstrate revolutionary combat capabilities, the Air Force can always procure additional platforms.

Finally, given the significant technological uncertainties and associated operational risks surrounding tactical autonomy, the Air Force must adopt and sustain a robust hedging strategy. This does not mean abandoning fundamental autonomy research but rather balancing investments across a portfolio of capabilities to manage risk. Priority investments should include advanced datalinks communication systems, enhanced human-machine interfaces, diverse methods for short- and long-range control of robotic aircraft, development of future weapons, and continued acquisition of advanced fighter platforms with enhanced sensors, processing capabilities, and low-observable technologies. By prioritizing these areas, the Air Force maximizes operational flexibility and avoids strategic vulnerabilities if autonomy fails to deliver promised capabilities.

Future methodical RDT&E investments may produce evidence that points to effective use cases for tactical autonomy within a combat environment. But technological realism demands that military leaders acknowledge present-day limitations and pursue advancements without relying solely on optimistic expectations.

Conclusion

The allure of AI-powered autonomous combat aircraft is powerful. But this vision of a future cost-effective force of tireless, precise machines unconstrained by human limitations collides with three fundamental problems that cannot be wished or engineered away: the challenge of perception in complex combat environments, the scarcity of relevant training data, and the unmitigated vulnerability of these systems to adversarial attacks. These cascading problems create limitations that incremental steps alone, such as tethering or limited CCA employment, cannot fully overcome.

The Air Force must immediately pivot to technological realism and account for these realities. While limited experiments with human-machine teaming may yield valuable insights, they do not offer a viable path to meaningful combat capability. Scaling even this reduced concept to future force structure planning without addressing the fundamental technical challenges inherent in tactical autonomy risks strategic failure. Institutional momentum behind overly ambitious autonomy planning risks not only financial misallocation but strategic inflexibility, limiting the Air Force's future options.

The path forward demands rigorous acquisition milestones tied to demonstrated capabilities, disciplined procurement aligned with technological maturity, and strategic hedging via balanced investment across a portfolio of proven technologies. As a warfighting organization, the Air Force's primary mission is to prepare for and conduct operations effectively rather than focusing on developing unproven technologies. Unless the Air Force implements standards for its 2030s force structure driven by evidence rather than optimism, it commits to a path that wagers US military defense on conjecture. Æ