[The US Army War College Quarterly: Parameters](https://press.armywarcollege.edu/parameters)

Volume 53
Number 1 *Spring Issue*

Article 14

3-3-2023

Minotaurs, Not Centaurs: The Future of Manned-Unmanned Teaming

Robert J. Sparrow

Adam Henschke

Recommended Citation

Robert J. Sparrow & Adam Henschke, "Minotaurs, Not Centaurs: The Future of Manned-Unmanned Teaming," *Parameters* 53, no. 1 (2023), doi:10.55540/0031-1723.3207.

Minotaurs, Not Centaurs: The Future of Manned-Unmanned Teaming

Robert J. Sparrow and Adam Henschke

ABSTRACT: Contesting Paul Scharre's influential vision of "centaur warfighting" and the idea that autonomous weapon systems will replace human warfighters, this article proposes that the manned-unmanned teams of the future are more likely to be *minotaurs*, teams of humans under the control, supervision, or command of artificial intelligence. It examines the likely composition of the future force and prompts a necessary conversation about the ethical issues raised by minotaur warfighting.

Keywords: manned-unmanned teaming, centaur warfighting, autonomous weapon systems, future force, ethics

What role will human beings play in the wars of the future? An influential answer to this question is that they will partner with sophisticated machines to leverage the distinctive capacities of both parties. Paul Scharre coined the term *centaur warfighting* to describe the use of manned-unmanned teams, arguing they possess a number of key advantages relative to the use of autonomous weapon systems (AWS).[1] By enabling human beings to control, supervise, or command multiple unmanned systems, human judgment and cognitive flexibility can be combined with the reaction speed, sensors, strength, and power of machines to outperform humans and machines fighting separately.

A centaur is a mythical creature with the head and upper body of a man and the lower body of a horse. When used to describe manned-unmanned teams, the image of the centaur promotes the idea that human beings will lead the team. We outline an alternative vision of the nature of manned-unmanned teams, which is more likely to be realized

Acknowledgments: Robert Sparrow is an associate investigator in the Australian Research Council Centre of Excellence for Automated Decision-making and Society (CE200100005) and contributed to this paper in that role. Adam Henschke's contribution was supported by the research program Ethics of Socially Disruptive Technologies, which is funded through the gravitation program of the Dutch Ministry of Education, Culture, and Science and the Netherlands Organization for Scientific Research (grant number 024.004.031). The authors would like to thank Tom Drummond for comments and discussion that have improved the paper. Joshua Hatherley assisted with bibliographic research for the paper.

1. Paul Scharre, "Centaur Warfighting: The False Choice of Humans vs. Automation," *Temple International & Comparative Law Journal* 30 (2016): 151–65; and see also Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W. W. Norton & Company, 2018).

in key domains of warfighting in future wars. Rather than human beings directing multiple robots, we suspect artificial intelligences (AI) will direct the activities of multiple human beings. The cyborg soldier of the future is more likely to be a *minotaur*—a mythical creature with the body of a man and the head of a bull—than a centaur: they will have a monstrous head rather than a monstrous body.

The reasons why the minotaur is a better figure for thinking about the future of human participation in the prosecution of war relate to technological dynamics and ethical imperatives. Artificial intelligences are arguably already more capable of performing the cognitive tasks most relevant to warfighting than robots are capable of performing the functions of the human body most relevant to warfighting. Moreover, advances in the applications of AI are emerging more rapidly than are advances in the applications of robotics. For the foreseeable future, then, in many domains, it will be more plausible to substitute machines for humans where humans have executive roles than where humans have roles involving the manipulation of objects or movement through cluttered environments. Indeed, there will often be an ethical imperative to place human beings under machines' control, supervision, or command. As the tempo of military operations increases due to the introduction of new technologies, shifting some functions of battlefield command to AI will help prevent friendly fire incidents and enhance the survivability of human warfighters.

Given the pace at which AI is being developed, there is an urgent need to consider the implications of *minotaur warfighting*, both for the effectiveness of the fighting forces of the future and for the human beings who will increasingly fight wars at the direction of machines. This article discusses the factors driving us toward a future in which wars are fought by minotaurs and begins a conversation about the ethical implications of minotaur warfighting.

Centaur Warfighting

Scharre introduced the idea of centaur warfighting in the context of debates about the impact and the ethics of the use of autonomous weapons. Scharre concedes that fully autonomous weapon systems may have a role to play in future warfighting, but he argues that, in most cases, teams of humans and machines will outperform both when they operate separately. Robots and AI programs excel at integrating large amounts of data, responding quickly, and carrying out precision strikes. At present, though, they are less capable than human beings when it comes to other roles critical to warfighting. In particular, he suggests, AI still struggles to make good decisions in complex and unexpected circumstances and, especially, to exercise the moral judgment necessary to resolve the ethical dilemmas that often arise in the context of war. For this reason, in many applications, according to Scharre, teams of human beings and robots (or AIs) working together will outperform autonomous systems and human beings when they fight separately, in combat and in other military operations.[2]

The idea of centaur warfighting has been highly influential, in terms of the way people understand the operations of existing weapon systems and as a model for how to design and use robots and AI to fight wars going forward.

Scharre himself uses the figure of the centaur to analyze the operations of the US counter-rocket, artillery, and mortar system (C-RAM), which has been in operation since 2010. This system consists of a radar-and-computer-controlled high-speed Gatling gun. It is highly automated and capable of engaging and destroying targets without human supervision. Current doctrine, however, requires a human to be in the loop (meaning the system cannot operate without input from a human operator) to authorize the engagement of particular targets to reduce the risk of fratricide. Scharre cites this arrangement as a model of centaur warfighting that should be emulated wherever possible. He also notes that similar systems include humans only "on the loop" (meaning the human supervisor has the option of intervening to alter the operations of the system) and suggests this arrangement may become increasingly common and even necessary as the tempo of operations increases due to the use of autonomous systems in more roles.[3]

Importantly, centaur warfighting serves as a model for future operations using robotic and autonomous systems currently under development. The US Department of Defense *Unmanned Systems Integrated Roadmap 2017-2042* asserts that "[m]ilitary operations of the future will require collaboration between unmanned systems and humans" and emphasizes the importance of "human-machine teaming." Manned-unmanned teaming (MUM-T) is a key goal of the *US Army Unmanned Aircraft Systems Roadmap 2010-2035* and also features heavily in the US Navy's *Unmanned Campaign*

2. Scharre, "Centaur Warfighting."
3. Scharre, "Centaur Warfighting," 157–60; and Scharre, *Army of None*, 323–30.

Framework.[4] While manned-unmanned teaming is compatible with a range of relationships between the machines and the humans in the team, the context and examples provided in document make it clear that these manned-unmanned teams are imagined as centaurs. Moreover, the flagship examples of unmanned systems include the US Army's Robotic Combat Vehicle-Light, the US Marine Corps' Remotely Operated Ground Unit for Expeditionary (ROGUE)-Fires platform, the US Air Force's SkyBorg project, and the Royal Australian Air Forces Loyal Wingman project (recently renamed MQ-28A Ghost Bat). These examples are almost universally advertised as enhancing the effectiveness of human warfighters—that is, as facilitating centaur warfighting.[5]

Human Beings in Charge?

The figure of the centaur implies that when humans cooperate with robots or AI, humans will be in charge of the team. Precisely how they will be in charge remains open. Scharre introduced the idea of centaur warfighting by using examples where human beings are *in* the loop and *on* the loop. Scharre seems to allow that centaur warfighting is compatible with cases where human beings play only a more distant supervisory role, though presumably, machines will need to remain under "meaningful human control" to speak of human supervision at all in this context.[6] Scharre's is a generous interpretation of what it means for the human being in a manned-unmanned team to be in charge. Even on this account, we believe, *contra* Scharre, that many collaborations of humans and machines will be more accurately described as minotaurs

4. Office of the Assistant Secretary of Defense for Acquisition, *Unmanned Systems Integrated Roadmap 2017–2042* (Office of the Secretary of Defense, 2018), 29, https://apps.dtic.mil/sti/pdfs/AD1059546.pdf; US Army UAS Center of Excellence, *US Army Unmanned Aircraft Systems Roadmap 2010–2035* (Fort Rucker, AL: US Army UAS Center of Excellence, 2010), 15–16, https://irp.fas.org/program/collect/uas-army.pdf; and Department of the Navy, *Unmanned Campaign Framework* (Washington, DC: Department of the Navy, 2021), 8, 11, https://www.navy.mil/Portals/1/Strategic/20210315%20Unmanned%20Campaign_Final_LowRes.pdf?ver=LtCZ-BPlWki6vCBTdgtDMA%3D%3D.

5. Eric Tegler, "An Army General Says the Robotic Combat Vehicles It's Experimenting With Will Be the 'Ghosts of Patton's Army,'" *Forbes* (website), July 30, 2021, https://www.forbes.com/sites/erictegler/2021/07/30/an-army-general-says-the-robotic-combat-vehicles-its-experimenting-with-will-be-the-ghosts-of-pattons-army/?sh=5c1c2d1666c8; Brian O'Rourke, "Rogue Fires," *Proceedings of the United States Naval Institute* 147, no. 6 (June 2021), https://www.usni.org/magazines/proceedings/2021/june/rogue-fires; Stephen Losey, "Air Force Aims to Sharpen Vision for Teaming Pilots with Drones," DefenseNews (website), March 18, 2022, https://www.defensenews.com/air/2022/03/18/air-force-aims-to-sharpen-vision-for-teaming-pilots-with-drones/; and Melanie de Git, "Loyal Wingman Uncrewed Aircraft Completes First Flight," Boeing (website), April 12, 2021, https://www.boeing.com/features/innovation-quarterly/2021/04/loyal-wingman.page.

6. Heather M. Roff and Richard Moyes, "Meaningful Human Control, Artificial Intelligence and Autonomous Weapons" (briefing paper prepared for the Informal Meeting of Experts on Lethal Autonomous Weapons Systems, UN Convention on Certain Conventional Weapons, April 2016), https://article36.org/wp-content/uploads/2016/04/MHC-AI-and-AWS-FINAL.pdf.

(unmanned-manned teams) than centaurs (manned-unmanned teams) because the machines in the teams will effectively be in charge.

Key Technological Dynamics

The initial period of the development of AI was characterized by the belief that the key challenge was to create machines able to perform cognitive tasks—playing chess, completing mathematical operations, and dealing with large datasets, for example—we find hard and think of as the pinnacle of intellectual achievement. For instance, notoriously, the original grant application for the Dartmouth Summer Workshop, which is widely recognized as the starting point of modern research into artificial intelligence, suggested that:

> An attempt will be made to find how to make machines that use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.[7]

It soon became apparent, however, that the real challenge lay elsewhere. Tasks we find easy (and that, we think, do not require intelligence because they are performed equally well by children and animals as by adults) proved to be difficult for machines. Perception, locomotion, and manipulation are now recognized as hard problems in AI and robotics. Despite significant progress in machine vision over the last decade, the capacity of machines to orient themselves and recognize objects in unstructured environments remains limited, and dexterous manipulation remains a key challenge. Robustness in real-world operating conditions is also challenging for robotic systems, as are energy requirements. It is striking at the current moment how much faster AI research is progressing than robotics research. In general, AI heads are still better than robot bodies. As computer scientist Donald Knuth observed, "AI has by now succeeded in doing essentially everything that requires 'thinking' but has failed to do most of what people and animals do 'without thinking'—that, somehow, is much harder."[8] Given the extent

7. As quoted in Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford, UK: Oxford University Press, 2014), 5.
8. Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (New York: Cambridge University Press, 2010), 318, quoted in Bostrom, *Superintelligence*, 14.

of the challenges involved in embodied engagements with the physical world, this will likely remain the case for the foreseeable future.

The relative strengths of AI and robotics are evident in many civilian applications. Today, when humans and intelligent machines work together in teams in industry, machines often perform the mental work while humans do the physical work.

Perhaps the most prominent civilian example of a minotaur is Amazon's fulfillment centers. Workers in these centers are directed and supervised by machines. Complex algorithms determine which goods must be shipped, where, and how, but humans must collect and (sometimes) pack them. Machines tell workers, via handheld devices, what to collect and from where. Because the warehouses use an algorithmic packing system, which stores goods to minimize the time required to collect them, rather than in a fixed location, humans could not find the products they are required to collect without the machines.[9] More and more, roboticized forklifts, pallets, or storage units bring product bins to tables where humans lift and package the goods or put them in another machine: human beings are thus reduced to being the hands of machines. As Noam Scheiber noted in the *New York Times*, "[the] steady stripping of human judgment from work is one of the most widespread consequences of automation—not so much replacing people with robots as making them resemble robots."[10]

Another example of a civilian minotaur is provided by the evolution of long-haul trucking. Truck drivers increasingly collect and deliver items and follow routes assigned to them by algorithmic logistics systems. Sensors that transmit data back to these algorithms monitor the drivers' speed, route, and driving performance. Drivers may even be automatically penalized for various infractions.[11] Instead of humans choosing a destination to which an autonomous vehicle drives, autonomous systems instruct humans when, where, and how to drive. These human-machine teams have developed and flourished, mostly because they are more technologically feasible than the

9. Chris Baraniuk, "How Algorithms Run Amazon's Warehouses," BBC Future (website), August 18, 2015, https://www.bbc.com/future/article/20150818-how-algorithms-run-amazons-warehouses.

10. Noam Scheiber, "Inside an Amazon Warehouse, Robots' Ways Rub Off on Humans," *New York Times* (website), July 3, 2019, https://www.nytimes.com/2019/07/03/business/economy/amazon-warehouse-labor-robots.html.

11. Robin Kaiser-Schatzlein, "How Life as a Trucker Devolved into a Dystopian Nightmare," *New York Times* (website), March 15, 2022, https://www.nytimes.com/2022/03/15/opinion/truckers-surveillance.html?searchResultPosition=2. See also Karen E. C. Levy, "The Contexts of Control: Information, Power, and Truck-driving Work," *Information Society* 31, no. 2 (2015): 160–74.

alternatives (such as humans choosing routes or machines driving), but also because they reduce road deaths.

The difference in the performance of AI and robots has been under-recognized in discussions of the future of manned-unmanned teaming in war because the success stories for unmanned systems in military applications have involved machines operating in (or attacking targets in) the aerial domain. Problems related to perception, navigation, and locomotion are relatively tractable for unmanned aerial vehicles, and problems related to manipulation and dexterous handling do not arise. If one considers the performance of unmanned systems at tasks central to land warfare, the situation looks very different.

Urban environments, forests, mud, snow, ice, and sand are extremely challenging for robots. Moving safely through such terrain in wartime requires constant judgments about how objects and surfaces interact, the best route to choose, the goals of other friendly and enemy units, and the information available to other agents. while humans make these judgments intuitively and often unconsciously, this bodily and perceptual know-how is difficult to render in algorithms. That the unstructured and refractory nature of the physical environment poses profound challenges for robots is even more obvious when it comes to other activities that play a key role in land warfare. For instance, transporting and emplacing ordnance, setting up defensive fortifications, or clearing a building requires humans. Increasingly, though, machines can identify enemy military objects and personnel in (near) real-time by integrating information from multiple sources (such as drones, satellites, video feeds from cameras mounted on weapons or helmets, and signals intelligence), a task that can exceed the capacities of humans, but is now within the capabilities of machines. Scharre and others have argued that machines may struggle to take account of contextual cues important to applying the law of war or make the ethical or strategic judgments required to determine an appropriate target.[12] In many circumstances, however, the context will actually make such decisions tractable for machines. For instance, it is plausible that in a particular area or engagement, all enemy submarines, tanks,

12. Scharre, "Centaur Warfighting"; Heather M. Roff, "The Strategic Robot Problem: Lethal Autonomous Weapons in War," *Journal of Military Ethics* 13, no. 3 (2014): 211–27; and Robert Sparrow, "Twenty Seconds to Comply: Autonomous Weapon Systems and the Recognition of Surrender," *International Law Studies* 91 (2015): 699–728.

or fighter aircraft may be legitimate targets, and it is within the capacities of existing AI to distinguish such systems from civilian objects.[13]

Similar observations can be made regarding the applications of AI and robotics in naval warfare. Many tasks essential to the operations of ships will be hard to automate or assign to robots because they rely on humans' ability to identify, move, and manipulate a range of different objects in complex environments. In particular, for the foreseeable future, humans will be needed to load and maintain weapons and service engines. However, cognitive tasks central to naval warfare (such as determining the best routes for ships, controlling air defense systems, and identifying and prioritizing targets) appear well within the capabilities of existing—or near-future—AI capacities.

Shifting our attention from the skies to land and naval warfare highlights the ways in which the physical environment poses challenges to the operations of machines that we are further from solving than we are from creating AIs that can identify targets and set priorities for warfighters.

Minotaur Warfighting

The emergence of minotaur teams in civilian life suggests minotaurs will also play a role in military operations in the future. At the very least, military stores and logistics will likely follow civilian models and create minotaur teams to carry out key functions.

Turning to combat operations, plausible use cases for minotaurs abound.

Despite recent progress in automating some functions, towing or driving, emplacing, and loading ordnance requires multiple humans. Identifying, tracking, and prioritizing targets, though, can be done by machines. In the era of so-called network-centric warfare, the best way to get inside the enemy's Observe-Orient-Decide-Act loop is to allow a computer to allocate targets or even aim and fire weapons.[14] Indeed, the Ukrainian military has reportedly taken significant steps in this direction in the current war with Russia.[15]

13. Robert Sparrow, "Robots and Respect: Assessing the Case Against Autonomous Weapon Systems," *Ethics and International Affairs* 30, no. 1 (Spring 2016): 93–116; Marcello Guarini and Paul Bello, "Robotic Warfare: Some Challenges in Moving from Noncivilian to Civilian Theaters," in Patrick Lin, Keith Abney, and George A. Bekey, eds., *Robot Ethics: The Ethical and Social Implications of Robotics* (Cambridge, MA: MIT Press, 2012), 129–44.

14. Wes Haga and Courtney Crosby, "AI's Power to Transform Command and Control," National Defense (website), November 13, 2020, https://www.nationaldefensemagazine.org/articles/2020/11/13/ais-power-to-transform-command-and-control.

15. See Charlie Parker, "Uber-style Technology Helped Ukraine to Destroy Russian Battalion," *Times* (website), May 14, 2022, https://www.thetimes.co.uk/article/uk-assisted-uber-style-technology-helped-ukraine-to-destroy-russian-battalion-5pxnh6m9p.

Where weapons lugged and loaded by human beings are aimed and fired at targets chosen by machines, we have minotaur warfighting.

Emerging technologies also threaten to turn infantry squads into minotaurs. The US Army's Integrated Visual Augmentation System will provide warfighters with tactical data using a mixed-reality headset based on Microsoft's HoloLens.[16] A recent US Army "request for information" provides clues as to how its developers anticipate AI will be used to extend the capabilities of this system.[17] This document lists "AI-enabled target detection algorithms," "machine assisted mission planning," "AI tactical predictions," and an "AI-enabled digital battlefield assistant" as areas of interest. Although the request's phrasing implies that AI will act as an assistant or adviser, there are strong reasons to believe that AI will not remain confined to these roles for long.

Studies of Human-Computer Interaction (HCI) show that people tend to over trust artificial intelligence, especially if the AI has proven itself generally reliable—a phenomenon known as *automation bias*.[18] If a target detection algorithm or battlefield assistant indicates that a particular object or person is a threat, warfighters are unlikely to gainsay the AI, especially given that the premise of the Integrated Visual Augmentation System is that it helps reduce the fog of war. Moreover, where the AI can draw on information from multiple platforms and sensors to formulate threat assessments or mission objectives, it may be wrong to act against its advice, given the machine's better vantage point. Once the performance of AI reaches a certain level, warfighters who assert their judgment over the AI's judgment will place their lives and the lives of those around them at risk; they will also detract from the combat effectiveness of the team. Eventually, the advice of AI will come to have the psychological, or even normative and institutional, force of orders, and warfighters engaged in small-unit combat will spend most of their time trying to achieve goals set for them by an AI.

As Thomas Adams argued more than two decades ago, as the impact of AI accelerates the tempo of battle and reduces effective decision-making time for humans, militaries may have little alternative but to outsource

16. "Next-generation Headset Preps Soldiers for Future Battlefield," US Army (website), November 13, 2020, https://www.army.mil/article/240851/next_generation_headset_preps_soldiers_for_future_battlefield.
17. US Army Contracting Command, Aberdeen Proving Ground, "Artificial Intelligence/Machine Learning Software Development and Integration," SAM.gov (website), January 19, 2022, https://sam.gov/opp/7201d5f4370d491e8322084cf58ec4e5/view.
18. Linda J. Skitka, Kathleen L. Mosier, and Mark Burdick, "Does Automation Bias Decision-making?" *International Journal of Human-Computer Studies* 51, no. 5 (1999): 991–1006, DOI:10.1006/ijhc.1999.0252.

many decisions to AI.[19] For the foreseeable future, though, the successful prosecution of war will involve human beings dealing with the mundane physical and material challenges machines currently handle poorly. If, in the future, AI is doing the cognitive work in a manned-unmanned (or, more accurately, unmanned-manned) team by choosing targets and setting goals, and humans are toiling at the direction of AI, we will have a minotaur rather than a centaur. The same dynamic Adams identified suggests that minotaurs will triumph over centaurs in future battles, creating a strong incentive for militaries to adopt minotaur warfighting.

The ultimate minotaur fighting force would consist of a team of humans and robots commanded by the AI equivalent of a general officer. Although not yet feasible, in the long term this idea is less far-fetched than it appears. Artificial intelligence tends to excel at games, including wargames, because an AI can learn from the experience of playing multiple iterations of a game.[20] If war were only moving units on a screen or maximizing a score according to a complex set of rules, machines would already outperform humans at directing military operations. The reason they do not yet relates to the difficulties involved in accurately representing in military simulations the capacities of different weapon systems and military units and the affordances of the terrain (including human terrain) on which operations will take place. Should the technology of military simulations improve so real-world operations can be accurately represented in wargames, the door will be opened for the development of sophisticated war-fighting algorithms.[21] Eventually, the pursuit of victory may require handing over command to machines and victory may be determined by which force has the better AI.

It is also worth noting that automation bias suggests some highly automated systems people currently think of as centaurs are actually minotaurs. If the human "in the loop" is unlikely to gainsay the machine, then the manned-unmanned team is a minotaur rather than a centaur. We suspect this may be the case with the Phalanx close-in weapon system and counter-rocket, artillery, and mortar system.

Finally, recognizing that even purportedly "autonomous" systems will rely on humans to load, repair, and maintain them suggests that

19. Thomas K. Adams, "Future Warfare and the Decline of Human Decision-making," *Parameters* 31, no. 4 (Winter 2001): 57–71, DOI:10.55540/0031-1723.2058.

20. Paul K. Davis and Paul Bracken, "Artificial Intelligence for Wargaming and Modeling," *Journal of Defense Modeling and Simulation* (February 2022), https://journals.sagepub.com/doi/abs/10.1177/15485129211073126.

21. Vinicius G. Goecks et al., "On Games and Simulators as a Platform for Development of Artificial Intelligence for Command and Control," *Journal of Defense Modeling and Simulation* (March 2022), https://journals.sagepub.com/doi/10.1177/15485129221083278.

many autonomous weapons systems should be understood as the heads of minotaurs and their human support teams being the body.

Ethical Implications

It is vital that military policymakers and the broader society begin a conversation now about the ethics of minotaur warfighting to prepare for or shape the future.

There are powerful ethical arguments *for* minotaur warfighting. Minotaur warfighting is likely to emerge in response to the ethical imperative to avoid fratricide. More controversially, the obligation of civilian society to warfighters, and of commanders to their troops, to avoid exposing friendly forces to unnecessary risk will also often argue in favor of minotaur warfighting. By swiftly identifying and prioritizing targets, minotaurs will reduce opportunities for the enemy to bring weapons to bear. Finally, the fact that minotaurs are likely to defeat centaurs in the not-too-distant future is ethically salient; if we fight in a just cause there are strong ethical reasons to field as powerful a military force as possible.[22]

Nevertheless, minotaur warfighting also has some profoundly troubling aspects. Indeed, in granting machines power over humans to the point of sending them into battle to be killed, minotaur warfighting foregrounds ethical questions being discussed in the contemporary debate about the relationship between machines and humans more generally.

One worry is that machines will not care sufficiently, or in the right way or, indeed, maybe at all about the lives of those they command. For instance, AI generals might use humans as cannon fodder to clear the way for a more powerful unmanned system. It is important to distinguish here between a worry that machines will risk human lives unnecessarily and a concern about them risking human lives at all. The former is really a doubt about the effectiveness of military AI and should eventually be assuaged by evidence that minotaurs win battles and reduce the risk to (human) warfighters.

The concern that human lives should not be at stake in the decisions of machines, which also arises in the debate about the ethics of using AWS, might be expressed in the philosophy of Immanuel Kant. Kant insisted that human beings should always be treated as "ends,"

22. Bradley Jay Strawser, "Moral Predators: The Duty to Employ Uninhabited Aerial Vehicles," *Journal of Military Ethics* 9, no. 4 (2010): 342–68.

not solely as means.[23] Unlike machines, humans have free will. According to Kant, we must respect this capacity in each other and avoid treating other people solely as tools to advance our purposes. It is difficult to see how machines could demonstrate such respect and easy to worry that minotaur warfighting could reduce human beings to mere means.

There is also a republican version of this objection. According to this tradition, liberty is compatible with laws that are the outcome of a process of deliberation that tracks the interests of citizens.[24] Where individuals can act as they wish only at the sufferance of the powerful, though, they are dominated and, to that extent, not free. The equal freedom of citizens requires that they not be subject to the arbitrary power of the sovereign or other citizens. It is tempting to think that the exercise of power by machines is *always* arbitrary insofar as machines cannot participate in the practices of reason giving that are constitutive of deliberation.[25]

Both these objections have merit. It is, however, difficult to formulate them in a way that does not invite the reply that a similar situation exists when humans order other humans into battle. Soldiers consent to be used to serve larger purposes when they enlist and are, arguably, subject to the arbitrary power of their superiors. While one hopes commanders only treat those under their command in ways they could justify to their subordinates in terms of their interests, military necessity may sometimes require otherwise. Thus, the ethics of command by machines does not look that distinct from the ethics of command more generally.

Nevertheless, it is hard to avoid a sense that there would be something wrong with granting machines the authority to send humans to their deaths. Humans are valuable in a way that machines are not. Placing humans under the command of machines seems to express the idea that machines are more important—or at least better—than humans. Unsurprisingly, this intuition also arises in the debate about the ethics of the use of autonomous weapon systems, wherein it plays an important role.[26]

Another question, which arises for both autonomous weapon systems and minotaurs, concerns the attribution of responsibility for decisions made by machines. When parents learn their children have been killed after being sent into battle by a machine, they may want to know whom to blame.

23. Immanuel Kant, *The Metaphysics of Morals* (New York: Cambridge University Press, 1991), 168.
24. Philip Pettit, *Republicanism: A Theory of Freedom and Government* (Oxford, UK: Clarendon Press, 1997).
25. Robert B. Brandom, *Reason in Philosophy: Animating Ideas* (Cambridge, MA: Harvard University Press, 2009).
26. Sparrow, "Robots and Respect."

When, if ever, it might be appropriate to hold a machine morally responsible for "its" actions remains a topic of vigorous philosophical debate.[27] We suspect the answer to the grieving parents' question will ultimately be settled as a matter of law, if not of morality, by assigning responsibility for the consequences of the decisions of AI and the actions of minotaurs to a human further up the chain of command.[28]

The ethical issues raised by minotaur warfighting identified here are troubling enough and more will undoubtedly emerge as the use of minotaur teams spreads. In military contexts, however, the case for minotaurs—that they will win battles and save the lives of friendly forces—is remarkably strong. For this reason, our analysis suggests that, as is often the case with new technologies, an all-things-considered ethical assessment of minotaur warfighting would require resolving a clash between Kantian and consequentialist intuitions.

Like many distinctions we use to understand the world, the contrast between centaurs and minotaurs is undoubtedly overdrawn. In reality, there will be a range of relationships between humans and machines when they work together. Even within particular teams, some tasks will be delegated more to humans and others to robots or to AI. Nevertheless, the image of the minotaur reminds us that this negotiation will not always favor human beings.

Similarly, insofar as the military is a system of systems, whether a particular collaboration between humans and machines is a centaur or a minotaur will be a function of the level of analysis. A system that looks like a minotaur if one draws the boundary around the team one way may appear as a centaur if one draws it another way. Thus, if AI battlefield assistants evolve to become AI squad leaders, one would hope human officers would command them. If the general directing the nation's military is an AI, one presumes the nation's civilian leadership would set the AI's war objectives. Recognizing, however, that, at some levels of analysis, machines will be in charge helps us understand the strengths and limitations of different forms of manned-unmanned teaming. As we have argued, the ethical questions will also look very different when one concedes that

27. Robert Sparrow, "Killer Robots," *Journal of Applied Philosophy* 24, no. 1 (2007): 62–77.
28. Peter Margulies, "Making Autonomous Weapons Accountable: Command Responsibility for Computer-guided Lethal Force in Armed Conflicts," in Jens David Ohlin, ed., *Research Handbook on Remote Warfare* (Cheltenham, UK: Edward Elgar Publishing, 2017).

some humans are effectively under the command of machines, even if those machines are in turn under the command of human beings.

Final Reflections

Minotaur warfighting will develop in key domains of warfighting because unmanned-manned teams will outperform manned-unmanned teams or humans or autonomous weapon systems operating alone. The nature of the relationships between humans and machines in this emerging mode of manned-unmanned teaming raises profound ethical questions. We also expect the development of minotaur warfighting to be challenging personally and institutionally for those who have spent their lives honing the capacity for human judgment that currently plays a central role in warfighting. Despite the ethical and institutional challenges posed by minotaurs, it is far from clear that those responsible for winning wars should resist the development of minotaur warfighting.

If militaries or societies do decide that putting warfighters under the control, supervision, or command of machines is a step too far, we believe that three tasks—one technological, one ethical, and one political—would need to be confronted as a matter of urgency.

First, significant financial and intellectual resources must be dedicated to developing robots capable of functioning effectively for extended periods in unstructured environments. A challenge in this task will be to succeed without also making it more plausible to hand over key cognitive tasks involved in warfighting to AIs. There is a real danger that the software advances required for robots to cope with the uncertainties and complexities of the physical environment will only further empower AI to strategize and exercise operational control over military forces.

Second, the intuitions grounding an ethical and political commitment not to put humans under the command of machines must be clarified and strengthened. Doing so also risks implying that the development and application of autonomous weapon systems is more problematic than many military ethicists and policymakers acknowledge. If it is morally wrong to allow machines to tell humans what to do, it is difficult to see how it could be morally permissible to allow machines to kill people.

Third, the international community must consider whether an international legal prohibition on the use of minotaur teams in war or in certain roles in war is desirable—or even feasible. Each nation will also need to consider how it will respond if other nations start fielding minotaur teams in war.

Whether any of these challenges can be successfully met, or if they should even be attempted, remains unclear. We hope this discussion has demonstrated the importance of confronting these questions. Before human warfighters cede the field to minotaurs, we need to know that the price of their victory will not be our humanity.

Robert J. Sparrow

Robert J. Sparrow is a professor in the philosophy program and an associate investigator in the Australian Research Council Centre of Excellence for Automated Decision-making and Society (CE200100005) at Monash University, Australia, where he works on ethical issues raised by new technologies. He has served as a cochair of the Institute of Electrical and Electronics Engineers Technical Committee on Robot Ethics and was one of the founding members of the International Committee for Robot Arms Control.

Adam Henschke

Adam Henschke is an assistant professor in the philosophy section at the University of Twente, Netherlands. His research is concerned with ethics, technology, and national security, and he is interested in ethical issues having to do with information technologies and institutions, surveillance, cyber-physical systems, human military enhancement, and relations between social information technologies, political violence, and political extremism. He recently coedited the 2022 *Palgrave Handbook of National Security*.

Select Bibliography

Adams, Thomas K. "Future Warfare and the Decline of Human Decision-making." *Parameters* 31, no. 4 (Winter 2001). DOI:10.55540/0031-1723.2058.

Davis, Paul K., and Paul Bracken. "Artificial Intelligence for Wargaming and Modeling." *Journal of Defense Modeling and Simulation* (February 2022). DOI:10.1177/15485129211073126.

Goecks, Vinicius G. et al. "On Games and Simulators as a Platform for Development of Artificial Intelligence for Command and Control." *Journal of Defense Modeling and Simulation* (March 2022). https://journals.sagepub.com/doi/10.1177/15485129221083278.

Haga, Wes, and Courtney Crosby. "AI's Power to Transform Command and Control." National Defense (website). November 13, 2020. https://www.nationaldefensemagazine.org/articles/2020/11/13/ais-power-to-transform-command-and-control.

Scharre, Paul. "Centaur Warfighting: The False Choice of Humans vs. Automation." *Temple International & Comparative Law Journal* 30 (2016).

Scharre, Paul. *Army of None: Autonomous Weapons and the Future of War*. W. W. Norton & Company, 2018.

Scheiber, Noam. "Inside an Amazon Warehouse, Robots' Ways Rub Off on Humans." *New York Times* (website). July 3, 2019. https://www.nytimes.com/2019/07/03/business/economy/amazon-warehouse-labor-robots.html?searchResultPosition=28.

Sparrow, Robert. "Robots and Respect: Assessing the Case against Autonomous Weapon Systems." *Ethics and International Affairs* 30, no. 1 (Spring 2016).