Cambridge Forum on AI: Law and Governance (2025), **1**, e48, 1–13
doi:10.1017/cfl.2025.10022

Cambridge Forum

RESEARCH ARTICLE

Augmenting military decision making with artificial intelligence

Karina Vold[1,2,3]

[1]Institute for the History and Philosophy of Science and Technology, University of Toronto, Toronto, ON, Canada; [2]Schwartz Reisman Institute for Technology and Society, University of Toronto, Toronto, ON, Canada and [3]Leverhulme Centre for the Future of Intelligence, University of Cambridge, Cambridge, UK
Email: karina.vold@utoronto.ca

(Received 20 January 2025; revised 16 July 2025; accepted 2 September 2025)

Abstract

As artificial intelligence (AI) plays an increasing role in operations on battlefields, we should consider how it might also be used in the strategic decisions that happen before a military operation even occurs. One such critical decision that nations must make is whether to use armed force. There is often only a small group of political and military leaders involved in this decision-making process. Top military commanders typically play an important role in these deliberations around whether to use force. These commanders are relied upon for their expertise. They provide information and guidance about the military options available and the potential outcomes of those actions. This article asks two questions: (1) how do military commanders make these judgements? and (2) how might AI be used to assist them in their critical decision-making processes? To address the first, I draw on existing literature from psychology, philosophy, and military organizations themselves. To address the second, I explore how AI might augment the judgment and reasoning of commanders deliberating over the use of force. While there is already a robust body of work exploring the risks of using AI-driven decision-support systems, this article focuses on the opportunities, while keeping those risks firmly in view.

Keywords: resort to force; military commanders; decision making; intuition; AI decision support; artificial intelligence

1. Introduction[1]

Much of the discussion of artificial intelligence (AI) in military contexts centres on its use in the operations or conduct of war, such as lethal autonomous weapons on the battlefield or in cyberwarfare. While these are vital areas of research and debate, they often eclipse another important potential application of AI: its use in the "War Room" (Erskine & Miller, 2024a, pp.137–8). That is, beyond the battlefield, AI could play a role in strategic deliberations that precede armed conflict, including critical decisions about whether it is right to resort to armed force. The potential role of AI in this earlier stage of decision making deserves closer attention. Decisions to resort to force are typically made by a small, elite group within a nation's political and military leadership. In these deliberations, senior military commanders play a crucial advisory role. They are consulted for their expertise in assessing the adversary's capabilities and intentions, outlining available military options, and anticipating the

[1]This is one of fourteen articles published as part of the *Cambridge Forum on AI: Law and Governance* Special Issue, *AI and the Decision to Go to War*, guest edited by Toni Erskine and Steven E. Miller.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (http://creativecommons.org/licenses/by-nc-sa/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use.

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

likely consequences of different courses of action (COA). This article explores how AI might be used to enhance the cognitive capacities of military leaders, particularly their ability to make complex, high-stakes judgments.

There are various ways that AI could support military commanders in this capacity, but these do not come without risks (Erskine & Miller, 2024b). There are well-documented potential risks that can emerge from introducing AI as a decision-support system into any human decision-making process. These risks include (but are not limited to) concerns about automation bias, diminishing the autonomy of human decision makers, skill atrophy and moral atrophy, safety concerns that could stem from this atrophy, concerns about blurred lines around moral and legal responsibility, and concerns about a loss of human control over AI. These risks are so serious that the idea of introducing AI as a decision-support system in such a high-stakes arena as the War Room may seem like an overall bleak idea. Why should we even consider the use of AI in such a context? Is there anything to be gained? If not, the way forward seems clear: simply do not use any AI-driven decision-support systems in this (or perhaps any) high stakes context. In this article, I play devil's advocate: against the backdrop of these risks, I consider whether this proposed use of AI (as decision support for military commanders) could bring any benefits or opportunities whatsoever to resort-to-force decision making. To do this, I explore two questions. (1) First, how do military commanders make decisions? (2) Second, how can AI be used to assist them in their critical decision making? To answer the first question, I draw on existing literature from psychology, philosophy, and military organizations themselves. To answer the second question, I explore possible applications of AI decision-support systems in context. I conclude that there are several ways in which *non-autonomous* AI could be applied to support senior military commanders.

If my argument stands, then we will have identified some possible benefits for the use of AI-driven supports in the context of resort-to-force decision making. Of course, identifying some potential benefits of a new technological intervention is not enough to justify its introduction. Any potential benefits must be weighed against the risks. Analysing the trade-off between potential benefits and risks to reach a conclusion about whether and when AI should be used in this context will fall beyond the scope of this article. I merely aim to identify *some* potential benefits that AI-driven decision-support systems could have in resort-to-force decision making; and to do so while recognizing that a number of high-stakes risks have already been identified (Erskine & Miller, 2024b). After all, if absolutely no potential benefits can be identified in favour of the use of AI-driven supports in this context, but there are clear risks, then it seems that the path to follow is straightforward. So, while this article aims to highlight some of the potential opportunities for using AI to support military commanders in critical roles, I leave the work of weighing these potential benefits against the risks to the reader. It could turn out to be a consistent and reasonable position to both accept the arguments I make (identifying potential benefits for how AI could augment and improve resort-to-force decision making), while still maintaining that the risks of this new technological intervention are too high and, hence, that AI should not be used in this context.

2. Non-rational agents

Before we consider how AI might augment the judgments of military commanders, it is important to first understand how human decision making works and why it falls short of being consistently reliable, "rational," or "good" (in some sense). The idea that humans are "rational agents" is a cornerstone of rational choice theory, a dominant theory in economics that assumes human actors always make rational choices, in the sense that they perform a calculation based on the information that is available to them and make the decision that provides them with the highest amount of personal utility. Economists often take rational choice theory to be value-free in the sense that the theory prescribes no distinctly moral or political behaviour to individual economic actors. Instead, agents are thought to simply respond to their preferences and demands, with their moral and political commitments already built into whatever preferences they have (Kahneman, 2011; Tversky & Kahneman, 1974).

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

Despite the prominence of rational choice theory in economics, a great deal of psychological research has shown that humans are not rational agents in this sense. Instead, studies have shown that human decision making is plagued by many evolved cognitive biases which lead to systematic deviations from rational judgments (Holyoak & Morrison, 2005; Kahneman, Slovic & Tversky, 1982; Pohl, 2004; Tversky & Kahneman, 1974). These biases emerged over our long evolutionary history. They can lead our decision making astray in all sorts of ways – extrapolating information from the wrong sources (salience bias) (Becker, Skirzyński, van Opheusden & Lieder, 2022; Bordalo, Gennaioli & Shleifer, 2012, 2013; Tversky & Kahneman, 1981), seeking to confirm existing beliefs rather than confront them (confirmation bias) (Nickerson, 1998; Wason, 1968), or even failing to remember events in the way they really happened (memory bias) (Kensinger & Schacter, 2006; Loftus & Palmer, 1974; Roediger & McDermott, 1995). Over 200 different biases (including anchoring bias, bandwagon effect, and implicit associations) have been documented, some better known than others (Korteling & Toet, 2021). Together, these biases challenge the description of human behaviour that rational choice theory offers (Kahneman, 2011).

It is important to note that our biases are not all bad, however. They can serve as helpful shortcuts for fast and good decision making, especially considering some of the brute physical limitations of our cognitive systems. For example, we are limited by how much information our brains can store, and how quickly we can retrieve that stored information and process it for real-time decision making (Kahneman, 2011). If a predator is quickly coming at you, you do not have time to decide the *best* way to escape. So, you probably will not make the all-things-considered, most optimal decision, as rational choice theory claims that you will, but if you make a good enough decision, and you make it fast enough, you will survive another day. This is likely why these biases have evolved: they serve an evolutionary function, increasing our odds of survival by allowing us to decide and act quickly based on limited amounts of information (Kahneman, 2011).

What is important for the purposes of this article is both that our cognitive "shortcuts" or "biases" can lead to problems in how military commanders make decisions, but also that these biases can enable commanders to make relatively good decisions in time-constrained conditions (even if they are not "perfectly rational," in the narrow sense intended by rational choice theory). In other words, while commanders rely on these biases as cognitive shortcuts (as humans inevitably do, regardless of context) this also means that there can be predictable deviations from the most optimal choice. This suggests that there is a gap that can be filled – a gap between what commanders decide and what is most optimal, given perfect information and sufficient time to weigh all options. With an awareness of this gap, researchers have begun to think more about how AI-driven systems can support human decision making across a variety of domains (e.g., Becker et al., 2022; Choi, Kim, Kim & Kang, 2021; Echterhoff, Melkote, Kancherla & McAuley, 2024; Shin, Kim, van Opheusden & Griffiths, 2023). Still, there is not much specifically on the military context, the decision making of commanders, or resort-to-force decision making, which will be the focus of this article (some recent discussions include Garcia, 2023; International Committee of the Red Cross [ICRC], 2024; Johnson, 2022; Sparrow & Henschke, 2023; Vold, 2024; Wheeler & Holmes, 2024).

3. Non-autonomous AI

Autonomous AI systems are those capable of autonomous decision making or action, where the system can perform tasks entirely on its own, with little or no human input, or that could even replace a human decision maker in some contexts. Some examples of these include lethal autonomous weapons, or even robot vacuums, self-driving cars, and robotic delivery systems. In each of these cases, the AI system makes and directly implements its own decisions. But there are many applications of sophisticated machine learning techniques in systems that do not aim to replace human decision makers. This vision of computing has a long history, dating back at least as far as Douglas Engelbart's work in the 1960s (Engelbart, 1962). These non-autonomous AI systems have varying

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

degrees of autonomy but generally share important common objectives: they are not built to replace human decision makers, and they are not built to function entirely on their own. Instead, they are built to support (e.g., by providing evidence or predictions) and to interact with humans. These include systems such as risk-predictor algorithms in policing, language translator assistants and prompt-enabled generative systems, such as ChatGPT and CoPilot.

The distinction between autonomous and non-autonomous systems is not always clear-cut.

Autonomous and non-autonomous AI systems could be more like two opposite ends of a spectrum, with many varieties in between. But for the purposes of this article, I will use non- autonomous AI systems to refer to (1) systems built to support, not replace, human decision making, and (2) systems built to interact with humans. Now, one could imagine circumstances where an AI system is designed to be non-autonomous but deployed or used in a way that it was not intended to be. For example, if a non-autonomous system is used to replace a human decision maker, it will likely not perform well given this is not what it was designed for. Hence, the distinction is important to make because these systems need to be treated and deployed differently, and they raise different risks and ethical concerns (Hernández-Orallo & Vold, 2019; Vold, 2024). For the rest of this article, I will focus on non-autonomous AI systems: those built to support, not replace, human decision making.

4. Risks of AI-driven decision support

While the goal of this article is to articulate some potential areas of opportunity for how non-autonomous AI-driven decision could benefit resort-to-force decisions, I also want to keep the risks of these technological interventions firmly in view. Furthermore, it is worth stating explicitly that identifying potential opportunities is not the same as endorsing deployment. Any potential benefits must get weighed against the risks. It is not my goal in this article to provide a comprehensive overview of all the possible risks associated with AI interventions into resort-to-force decision making, nor is it my goal to analyse the trade-off between potential benefits and risks. For reasons of space, these tasks must fall beyond the scope of this article. Still, because the decision context under discussion has such moral gravity, I want to highlight *some* of the risks of using AI-driven decision-support systems.

Many believe that AI systems are "objective" and, hence, that they can enhance objectivity in human decision-making processes (see discussions in Benjamin, 2019; Cave, Nyrup, Vold & Weller, 2019; Corbett-Davies, Goel & Gonzalez-Bailon, 2017; Johnson, 2023). However, there is a great deal of evidence against this claim to objectivity which instead shows how algorithms embed biases, often stemming from skewed training data, flawed assumptions or proxy measures, opaque model design, or poor deployment choices. Algorithmic bias undermines the purported objectivity of system predictions and outputs, which often have less accurate performance for particular groups of people, including racial, ethnic and linguistic minorities (Broussard, 2024; Buolamwini & Gebru, 2018; Eubanks, 2018; Fazelpour & Danks, 2021; Johnson, 2021; Noble, 2018; Ntoutsi et al., 2020; Shin & Shin, 2023). One important implication of this for our purposes is that while AI-driven decision-support systems may help overcome some of our built-in cognitive biases, they can also introduce new biases and reinforce existing biases.

Another well-documented risk of AI-driven decision-support systems is known as "automation bias" (Bainbridge, 1983). This occurs when users defer to algorithmic outputs rather than applying their own judgment. For example, studies suggest that individuals and teams who frequently rely on AI-driven systems may begin to accept the outputs of these systems with little or no scrutiny (Alon-Barkat & Busuioc, 2023; Goddard, Roudsari & Wyatt, 2012). This growing reliance can be especially problematic when system-generated recommendations are flawed. So, for example, when automation bias and algorithmic bias intersect, users may uncritically accept recommendations that are systematically biased against certain populations. This can lead to harmful outcomes depending on the context.

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

A further downstream effect of increased use and reliance on AI tools for decision support and cognitive offloading is the degradation of a user's skills. For example, over-time human decision makers may exercise their own expertise and judgement less and less, which could lead to skill atrophy. Some skills that might dimmish over time include things like critical thinking, situational awareness, intelligent and expert decision making (e.g., Cummings, 2017; Gerlich, 2025; Skitka, Mosier & Burdick, 1999; Zerilli, Knott, Maclaurin & Gavaghan, 2019). This gradual atrophy could in turn reduce the ability of human users to recognize when AI systems are failing, or to be able to intervene and make decisions without them. These effects contribute to the risk of complacency errors: mistakes that are made when a human user becomes overly reliant, passive, or inattentive and fails to actively monitor the outputs of a system, including failing to detect problems, anomalies or malfunctions (Bahner, Hüper & Manzey, 2008).

These risks feed into other concerns. For example, the introduction of "intelligent" technologies can complicate chain-of-command decision making and lead to liability and accountability gaps (Erskine, 2024; Sienknecht, 2024). When an AI system contributes to a poor decision, for example, it can be difficult to assign responsibility for fault. This is especially true when the internal functioning of AI systems, like deep neural networks, are so complex that their "reasoning" becomes opaque, even to the system designers. This is increasingly the case, as machine learning models are now highly complex, relying on many hidden layers of processing between input and output. The lack of opacity of these models creates problems for trust (Kares, König, Bergs, Protozel & Langer, 2023) and "explainability": we cannot explain how their outputs are reached. In some contexts, we may not need an explanation of how an AI system reaches its predictions (e.g., in game playing), but in other more high-stakes contexts we might, especially in circumstances where we need to rely on them or may want to challenge them. For each of these risks, there have been mitigation strategies suggested. For example, some studies show that trust in an AI system is modulated by a variety of factors, including transparency in how the system functions and when/how the system is used or deployed (discussions in Kares et al., 2023; Zerilli, Bhatt & Weller, 2022; Zerilli et al., 2019). While some evidence suggests that training and educating users about how models function can mitigate the risk of automation bias and complacency errors (Goddard et al., 2012; Masalonis, 2003). Still further empirical research is needed on whether and how it may be possible to avoid and mitigate these risks.

5. The military commander's decision-making process

When a state deliberates on whether to resort to force, political leadership must rely on critical information and guidance from military leadership. How senior military commanders make decisions about what military options are available (what's possible, what's risky, what's realistic and so forth) or predictions about what the potential outcomes of military actions might be is therefore highly consequential for state decisions about when to use force. So how do military commanders come to form their judgments on such matters?

One recent study based on 30 interviews with former and current NATO commanders and senior staff officers concluded that military commanders make key decisions in the planning process based on professional judgment, discretion and intuition (Sjøgren, 2022). AI might be able to assist in augmenting each of these – judgement, discretion and intuition – but in this article there is only space to examine the case for one of these in-depth and so here I focus on the role of intuition in decision making. Recall that my goal in this article is only to identify *some* potential benefits that AI-driven decision-support systems could have in resort-to-force decision making; not to do identify *all* possible areas of application. The latter would be too comprehensive and would require a much lengthier treatment. Consider that beyond decision making there are many other relevant cognitive capacities that AI could augment. Hernández-Orallo and Vold (2019), for example, suggest how AI could enhance memory, sensory (e.g., visual and audio) processing, attention and search functions, planning, communication, navigation, conceptualization, quantitative and logical reasoning,

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

and metacognition. All of these are possibilities worthy of exploration. With that said, let us turn to our focus: the role of intuition.

What is intuition? Intuition is often seen as one of the hallmarks of expert decision making. For example, a radiologist can look at an X-ray and diagnose a disease in a moment, and a world champion of Go will report "seeing" his next move straight away (Gobet & Chassy, 2009). Two prominent and early philosophers of cognition and AI, Hubert Dreyfus and Herbert Simon, agreed that intuition is one of the defining features of expertise (Dreyfus, 1972; Simon, 1989), though they proposed different mechanisms for how it functions in the brain. They agreed that some of the key elements of intuition are its rapid onset, its perceptual nature, and its holistic nature – allowing experts to quickly visualize and recognize the relevant features of a complex situation in its entirety. Intuitive decision makers can make quality decisions faster because their expert training and knowledge enables them to quickly separate what factors are relevant or important to consider.

How do we develop intuition? Good intuition takes time to cultivate. It is domain-specific – one cannot be skilled in intuitive decision making across random environments or without training (Kahneman & Klein, 2009; Klein, 2007). Learning situational elements, that is, elements that depend on context, seems to come with the normal stages of expert knowledge development, which also corresponds with an increase in intuitive understanding and a decrease in analytic thinking about what to do next (Phillips, Klein & Sieck, 2004). It is thought that building intuition requires deliberate practice, an environment in which it is safe to fail, and a focus on whether and why a plan worked, rather than an analysis about whether a codified process was followed to the letter (Gobet & Chassy, 2009; Sjøgren, 2022).

Despite its relatively good reputation, some influential perspectives approach intuition with scepticism. Kahneman, known for his foundational work on cognitive biases, has argued that our intuitions are shaped by these very biases, making them vulnerable to systematic judgement errors – that is, departures from rational decision making. As a result, he and others suggest that intuitive judgments should be evaluated through structured procedures, or even algorithmic methods, to ensure greater reliability (Kahneman, 2011; Kahneman & Klein, 2009). Rather than adopting a stance of scepticism toward intuition, I argue that intuition is not only unavoidable but essential, particularly for humans faced with high-stakes decisions in complex and uncertain environments. And instead of using AI to evaluate or override human intuition, we should use it as a tool to augment intuition. The goals should not be to replace or limit intuition in expert decision making, but to sharpen it. In the next section, I expand on three ways in which non-autonomous AI systems might be developed to do this for military commanders: by limiting the number of COAs to develop, by assisting with mental modelling and visualization techniques, and by monitoring and identifying human factors that can impact decision making.

6. Using AI to support a commander's intuitive decision making

In this section, I discuss three aspects of the military commander's decision-making process that could benefit from an application of non-autonomous AI by helping to train commander intuition.

6.1. Limit the number of Courses of Action (COA) to develop

In 2023, the U.S. Center for Army Lessons Learned (CALL) released a handbook on the military decision-making process (MDMP) (Mueller & Kuczynski, 2023). In this handbook, they note the following:

> Limiting the number of COAs developed and analyzed can save planning time. If available planning time is extremely limited, the commander can direct development of only one COA. In this case, the goal is an acceptable COA that meets mission requirements in the time available. This technique saves the most time. The fastest way to develop a plan has the commander

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

> directing development of one COA with branches against the most likely enemy COA or most damaging civil situation or condition. However, this technique should be used only when time is severely limited. In such cases, this choice of COA is often intuitive, relying on the commander's experience and judgment. (p. 7)

As mentioned in this passage, the intuition of a military commander becomes especially critical when time is severely limited. A properly trained AI could aid in the process of limiting the number of COAs that are developed and analysed by military teams (commanders and staff). These AI systems could be especially useful in time-constrained decision-making contexts. Imagine a system that is trained to analyse a wide range of possible COAs within a particular mission context. An AI system could analyse situational characteristics of the area of operations, including the terrain, the weather, the status of enemy forces and friendly forces (including, for example, their available equipment, personnel, systems, etc.), civilian considerations (who may be affected, what possible support might be needed from local civil authorities, etc.), and other mission details. Already this number of variables is difficult for a human to consider under pressure and time constraints.

How feasible is it for an AI system to do this? Much of this data is already digitally available and easily accessible for an AI system to train on. AI systems can pull terrain data from satellite imagery and digital elevation models from open sources (like Goole Earth), for example. This information can then be used to infer elevation, slopes, vegetation coverage and accessibility. Information about the historical weather patterns, current and forecasted weather is all easily accessible in real-time from public APIs, satellite feeds and radar data. Demographic information is less readily available (and certainly harder to access in real-time), but it is not impossible to access. Much of this information is available through Google Earth, satellite images and government databases. Information of about one's own military capacities is likely to be highly accurate and easily accessible. Considering these variables along with a mission objective analysis, an appropriately trained AI system could detect complex non-linear relationships between input variables (available actions) to identify the most compatible COA matches that can achieve the objective. The outputs of the AI system might include a ranking of COAs using a decision matrix to aid a comparison process.

Importantly, even with a non-autonomous AI-support system like this in place, commanders would still have to make intuitive decisions about what COA to take. These support systems would not be replacing human judgment. Instead, the system would be assisting the commander's judgment by taking into consideration multiple complex variables and offering a ranking of possible COAs. In this case, we can view the system as providing another piece of evidence for a commander to take into consideration in their own decision-making process (one that we understand relies on intuition). The ideal outcome would be for their professional expertise, intuition and final judgments to be supported, not replaced, by the AI system.

6.2. Mental modelling and visualization

Another critical cognitive capacity for military commanders is visualization. In cognitive science, we often speak about "mental models." When you know how to navigate your neighbourhood to get to a new place without looking at a map, this is because you have a mental model of the neighbourhood. Mental models are internal representations of external reality. They are often compressed or distilled versions of how something works. It is hypothesized that these models play a critical role in cognition, reasoning and decision making. Mental models allow us to reason and solve problems, and thus facilitate internal visualizations. There is a rich philosophical literature on mental models and how they work, and even more literature simply on models and how models work. One key idea is that models (in general) leave non-essential information out – they simplify. We cannot carry around a perfectly detailed image of the world in our heads, so a mental model captures only select concepts

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

and relationships that we (somehow) deem to be the most relevant for representing the "core" of what we are modelling.

In general, models are thought to be useful precisely because they reduce information loads. They reduce complex scenarios to include only core objects and relations (think of Watson and Crick's model of the DNA double helix, for example). Mental models are seen as especially critical for our spatial reasoning capacities, as tools that allow us to arrive at solutions to cognitive problems (Hemforth & Konieczny, 2006; Jones, Ross, Lynam, Perez & Leitch, 2011).

Our mental models are valuable sources of information that inform our decision making. And military commanders, like all humans, rely critically on mental models in decision-making processes. This is captured in the following passage of the U.S. CALL handbook on MDMP:

> A commander's ability to understand, visualize, describe, and direct during the operations process is paramount. … Successful commanders visualize the overall fight and give specified visualization to each staff section across each [area] and phase of the operation. Commanders mentally develop and maintain a running estimate for effective visualization, planning, and guidance… (p. 11)

This passage illustrates how a commander's capacity to visualize a combat zone and maintain a mental model of the area of operation is critical for successful time-constrained decision making, which in turn relies on intuition. And while this passage is about operational decisions, the same holds for strategic decisions. Within war room deliberations about whether to use force, the role of commanders is to provide information about operational and tactical options and outcomes. This is another area where a carefully designed non-autonomous AI system can support military commanders. When deciding how to prioritize between different tactical manoeuvres, for example, an AI system could consider relevant factors, such as the task and purpose of manoeuvres, forms of manoeuvres, reserve composition, priorities and control measures, the need for reconnaissance and surveillance, possible tactical deception, collateral damage or civilian causalities, any manoeuvre that could risk the achievement of end-state objectives, etc.

How feasible is it for an AI system to do this? It would be relatively straightforward to train an AI system on military doctrines, field manuals and historical battlefield data (including evaluative data on successes and failures). From this a system could infer the patterns and objectives of common military manoeuvres, such as envelopments, feints, frontal assaults, retreats and so on. It could also infer possible patterns of enemy deception. A nation's military has access to data about its own internal reserve composition (and cases with historical analogues), which can be used to inform AI models as they rank different possible manoeuvres. Reconnaissance and surveillance needs could be inferred by analysing terrain (from satellite data), line-of-sight paths, drone coverage, and combining data from multiple types of sensors (e.g., cameras, radars, thermal, acoustic), for example. Using this kind of information, a well- trained model could optimize the prioritization process by uncovering non-linear and subtle correlations among the many relevant variables that cannot be identified by a single person alone or by using conventional analysis models. As described, this data is already available to military commanders (some of it is available publicly), so introducing an AI system here does not need to increase surveillance. The data input into such an AI system would also not require a huge amount of continuous work from human personnel to continue to function, since much of the data can be scraped from existing digital databases.

As a cautionary note here, it is worth emphasizing that current frontier AI systems are primarily predictive models, designed to predict the mostly likely scenarios based on patterns in the data they have been trained on. Their core function is statistical pattern recognition, which is to say that they do not predict the future but rather, they generate plausible continuations of scenarios (e.g., conversations or images) based on past data and learned correlations. Understanding the appropriate function and limitations of these models is critical for any human decision makers that are supported

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

by these systems. Evidence suggests that this training and transparency around how models function can help prevent some of the risks discussed earlier, such as automation bias, complacency error and accountability gaps (Bahner et al., 2008; Goddard et al., 2012; Kares et al., 2023; Masalonis, 2003; Skitka, Mosier & Burdick, 2000; Zerilli et al., 2022).

Our mental models are useful for thinking through hypothetical future scenarios, or even imaginary scenarios, as much as they are for real situations. This means that an AI system that can help a human agent make more accurate predictions about future scenarios can also augment their mental modelling capacities, making their decision making and reasoning (based on these models) more effective. There are different ways to envision how this could be implemented, but one possible method would be to have the AI assistant produce diagrams to assist with visualizations. Evidence suggests that learning mental model development is enhanced by use of diagrams designed to highlight important structural relations (Butcher, 2006). I want to emphasize that, as before, an AI system in this context is not making any final decisions about COAs. These support systems are not final decision makers. Instead, their output is a supporting tool meant to inform a military commander in their decision-making process, not to replace the commander.

6.3. Human factors

Finally, a third area in which an AI could help improve decision making regards so-called "human factors." Human factors include things like emotional, mental, and physical fatigue. For example, CALL handbook on MDMP notes, "Commanders usually begin to leave out important pieces of guidance that would enhance efficient planning actions the longer the operation continues (because of fatigue, sleep deprivation, and compressed timelines)" (p. 13). This article is focused on strategic decisions about when to wage war, not operational decisions that are made during the conduct of war, but human factors still apply. Consider how high emotions were on 12 September 2001 in the White House Situation Room. The top advisors to the U.S. President who crowded into the jammed war room were described as having "eyes puffy from lack of sleep" (Carney & Dickerson, 2001). These are often the conditions under which states must decide whether it is right to use armed force.

Let us return now to the earlier discussion about the role of cognitive biases in human decision making. We established that these biases may serve some purpose and can help us make good decisions under time constraints, for example, but can also lead us astray. Some make a similar argument about emotions, for example, feeling fear at the sight of a predator, or empathy when another human is injured. Contrary to what traditional rational choice theory would say, our emotions play an inevitable role in decision making (Kahneman, 2011). Emotions sometimes override logical reasoning, of course, but other times they explain what most people would describe as rational behaviour (De Sousa, 1990). Consider that many legal systems have a category mitigating circumstances that recognizes how intense emotions like rage or jealousy can contribute to bad decisions, like acting violently or even murder. These crimes of passion make sense to us, even if they lead to very bad decisions. What is important for our purposes is to understand how emotional stress and fatigue can (sometimes) negatively impact human decision making. Emotions can run high, causing overreactions, rashness and even violence. Emotional, mental and physical fatigue can all contribute to poor and erratic decision making (Brownlee & Ungerleider, 2019; Danziger, Levav & Avnaim-Pesso, 2011; Persson, Barrafrem, Meunier & Tinghög, 2019). This suggests there is a space for thinking about how emotional, mental and physical fatigue may affect resort-to-force decision making and whether AI could be used to assist and complement human skills (Wheeler & Holmes, 2024).

Providing the kinds of support mentioned earlier for mental modelling tasks and planning different COA could indirectly support commanders by making their cognitive work less intense, putting them in a better emotional and physical state to make good decisions rapidly. But we could also imagine AI systems explicitly designed to help monitor and reduce the negative impact of emotional and physical fatigue. Much of this technology is already available. Biometric wearable technologies (e.g., a

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

smart watch or ring) can track states that indicate emotional and physical fatigue, for example, when our heart rates are too high or when we have been sitting too long and nudge us take actions to reduce these states. These could also monitor indicators like delayed reaction times, oxygen levels, number of hours on the job, number of hours of sleep, and so forth, and use this information to inform commanders about their own internal states. People are sometimes not aware of their own levels of stress, anxiety or fatigue: we forget to stand up, move around, take breaks, go to sleep, eat, etc. AI tools, like wearables, can alert users to signs of stress or fatigue that they may not notice, while also promoting exercises and actions for users to take to reduce these states of fatigue. Beyond physical indicators, there are other kinds of data that can improve the accuracy of the emotional intelligence of AI (i.e., the ability to read the emotional state of human agents). For example, the language we use when speaking with others, when writing or when engaging with a text-based AI system also indicates our emotional states (Lee, Paas & Ahn, 2024), and this analysis could likewise inform us about our current states, thereby increasing self-awareness about when we are most capable of making clear-headed decisions (in the sense of being less influenced by emotional or physical fatigue).

Importantly, unlike the AI systems discussed in previous sections, these AI systems require a significant amount of personal data to be accurate and effective. If biometric tracking systems are not used voluntarily and transparently, they may feel invasive, coercive and like a form surveillance. Moreover, the effectiveness of these systems will depend on how well they are designed and used, so having user consent (and buy-in) is even more important. There are ways to design these systems such that personal data is better protected and is accessible only to the user (e.g., using on-device processing and encryption methods), but this requires careful product design (Prakasha & Sumalatha, 2025).

7. Conclusion

A nation's deliberation about when to use armed force usually involves only a small group of powerful people: political leaders, top advisors, and the most senior military commanders. This article has considered whether, despite serious and well-documented risks, there may be any potential benefits to incorporating AI decision-support systems into the high-stakes context of resort-to-force deliberations. To narrow this examination, I have focused on the role of military commanders in these deliberations. By examining how military commanders typically make decisions by relying on their expert intuition, and exploring how AI could be used to augment these processes, I have argued that certain forms of AI could plausibly enhance strategic reasoning in the War Room. To make this proposal as concrete as possible, I have focused on three aspects of human decision making that are especially relevant for military commanders. These are (1) limiting the number of COAs, (2) improving mental modelling and visualization capacities, and (3) limiting the negative effects of human factors, including emotional and physical fatigue. I argue that non-autonomous AI systems could aid COA development and prioritization, improve situational awareness, offer alternative simulations to challenge biases, focus cognitive effort, and reduce negative effects of fatigue.

These areas are surely not the only areas in which AI systems have the potential to augment resort-to-force decision making. I have narrowed my focus on these to provide a more concrete and thorough treatment, rather than a high-level overview of opportunities. Hernández-Orallo and Vold (2019) have suggested multiple other cognitive capacities, besides decision making, that could be targeted for enhancement by non-autonomous AI systems, and that likely play an important role in how military commanders and other top officials deliberate about when to use force. These capacities include things like memory processes, sensory (e.g., visual and audio) processing, attention and search functions, planning, communication, navigation, conceptualization, quantitative and logical reasoning, and metacognition. Future research could explore these possibilities in more depth.

It is worth ending on a cautionary note. While my goal has been to highlight the potential opportunities, this is not the same as endorsing deployment. As emphasized throughout, the risks of

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

using AI-driven support systems remain profound: from automation bias and moral deskilling to accountability gaps, surveillance, and erosion of human judgment. The case for or against using AI in this context ultimately hinges on how these risks are managed, and whether the benefits meaningfully improve decision making without compromising ethical and legal responsibility. This article stops short of endorsing implementation. Instead, my goal is to encourage a more balanced debate – one that neither embraces AI nor rejects its potential uncritically. If AI is to play any role in decisions of such moral gravity as resort-to-force decisions, that role must be designed with caution, transparency and humility.

Funding statement. This work is partially funded by an AI2050 Early Career Fellowship with the Schmidt Sciences Foundation.

References

Alon-Barkat, S., & Busuioc, M. (2023). Human-AI interactions in public sector decision making: "Automation bias" and "selective adherence" to algorithmic advice. *Journal of Public Administration Research and Theory*, *33*(1), 153–169. https://doi.org/10.1093/jopart/muac007

Bahner, J. E., Hüper, A.-D., & Manzey, D. (2008). Misuse of automated decision aids: Complacency, automation bias and the impact of training experience. *International Journal of Human-Computer Studies*, *66*(9), 688–699. https://doi.org/10.1016/j.ijhcs.2008.06.001

Bainbridge, L. (1983). Ironies of automation. *Automatica*, *19*(6), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8

Becker, F., Skirzyński, J., van Opheusden, B., & Lieder, F. (2022). Boosting human decision-making with AI-generated decision aids. *Computational Brain Behavior*, *5*(4), 467–490. https://doi.org/10.1007/s42113-022-00149-y

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim*. Cambridge, England; Medford, Massachusetts: Polity Press.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, *127*(3), 1243–1285. https://doi.org/10.1093/qje/qjs018

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, *121*(5), 803–843. https://doi.org/10.1086/673885

Broussard, M. (2024). *More than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*. Cambridge, Massachusetts, UK: The MIT Press.

Brownlee, J. R., & Ungerleider, J. D. (2019). Decision making and ethics. In R. M. Ungerleider, J. N. Meliones, K. N. McMillan, D. S. Cooper, & J. P. Jacobs (Eds.), *Critical heart disease in infants and children* (3rd ed., pp. 73–81e1). New York, NY, USA: Elsevier. https://doi.org/10.1016/B978-1-4557-0760-7.00009-7

Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional Accuracy disparities in commercial gender classification. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency, PMLR*, *81*, 77–91.

Butcher, K. (2006). Learning from text with diagrams: Promoting mental model development and inference generation. *Journal of Educational Psychology*, *98*(1), 182–197.

Carney, J., & Dickerson, J. F. (2001, December 31). *Inside the war room*. *Time Magazine*. https://content.time.com/time/magazine/article/0,9171,1001454,00.html. 3rd January 2025.

Cave, S., Nyrup, R., Vold, K., & Weller, A. (2019). Motivations and risks of machine ethics. *Proceedings of the IEEE: Machine Ethics: the Design and Governance of Ethical AI and Autonomous Systems*, *107*(3), 562–574.

Choi, S., Kim, N., Kim, J., & Kang, H. (2021). *How does AI improve human decision-making? Evidence from the AI-powered Go program*. USC Marshall School of Business Research Paper, Sponsored by iORB, Los Angeles, California. https://doi.org/10.2139/ssrn.3893835

Corbett-Davies, S., Goel, S., & Gonzalez-Bailon, S. (2017, December 20). Even Imperfect algorithms can improve the criminal justice system. *The New York Times*, *The Upshot*. https://www.nytimes.com/2017/12/20/upshot/algorithms-bail-criminal-justice-system.html. 3rd January 2025.

Cummings, M. L. (2017). Automation bias in intelligent time-critical decision support systems. In D. Harris and W.-C. Li (Eds.), *Decision making in aviation* (pp. 289–294). Surrey, UK: Routledge.

Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011). Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences*, *108*(17), 6889–6892. https://doi.org/10.1073/pnas.1018033108

De Sousa, R. (1990). *The Rationality of Emotions*. MIT Press: Cambridge, Massachusetts, UK.

Dreyfus, H. L. (1972). *What computers can't do: A critique of artificial reason*. New York, NY, USA: Harper & Row.

Echterhoff, J. M., Melkote, A., Kancherla, S., & McAuley, J. (2024). Avoiding decision fatigue with AI-Assisted decision-making. *Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization (Vol. 6*, pp. 1–11). New York, NY, USA: Association for Computing Machinery.

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework* (Summary Report AFOSR-3223). Stanford Research Institute.

Erskine, T. (2024). Before algorithmic Armageddon: Anticipating immediate risks to restraint when AI infiltrates decisions to wage war. *Australian Journal of International Affairs*, *78*(2), 175–190. https://doi.org/10.1080/10357718.2024.2345636

Erskine, T., & Miller, S. E. (2024a). AI and the decision to go to war: Future risks and opportunities. *Australian Journal of International Affairs*, *78*(2), 135–147. https://doi.org/10.1080/10357718.2024.2349598

Erskine, T., & Miller, S. E. (Eds.). (2024b). *Anticipating the future of war: AI, automated systems, and resort-to-force decision making*, Special Issue. *Australian Journal of International Affairs*, *78*(2), 135–147.

Eubanks, V. (2018). *Automating inequality: How high-tech tools*. In *Automating inequality: How high-tech tools profile, police, and punish the poor*. New York, NY: St. Martin's Press.

Fazelpour, S., & Danks, D. (2021). Algorithmic bias: Senses, sources, solutions. *Philosophy Compass*, *16*(8), e12760.

Garcia, D. (2023). Algorithms and decision-making in military Artificial Intelligence. *Global Society*, *38*(1), 24–33. https://doi.org/10.1080/13600826.2023.2273484

Gerlich, M. (2025). AI Tools in Society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, *15*(6), 14–28. https://doi.org/10.3390/soc15010006

Gobet, F., & Chassy, P. (2009). Expertise and intuition: A tale of three theories. *Minds & Machines*, *19*, 151–180. https://doi.org/10.1007/s11023-008-9131-5

Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation Bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, *19*(1), 121–127. https://doi.org/10.1136/amiajnl-2011-000089

Hemforth, B., & Konieczny, L. (2006). Language processing: Construction of mental models or more? In C. Held, M. Knauff & G. Vosgerau (Eds.), *Mental models and the mind* (pp. 189–204). North Holland, Netherlands: Elsevier.

Hernández-Orallo, J., & Vold, K. (2019). AI extenders: The ethical and societal implications of humans cognitively extended by AI. In V. Conitzer, G. Hadfield & S. Vallor (Eds.), *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 507–513). Association for Computing Machinery. https://doi.org/10.1145/3306618.3314238

Holyoak, K. J. and Morrison, R. G. (Eds.). (2005). The Cambridge. In *handbook of thinking and reasoning*. Cambridge, UK: Cambridge University Press. https://doi.org/10.1017/CBO9780511998005

International Committee of the Red Cross (ICRC). (2024). *International humanitarian law and the challenges of contemporary armed conflicts: Building a culture of compliance for IHL to protect humanity in today's and future conflicts*. ICRC Humanitarian Law and Policy Blog. Ref. 4810.https://www.icrc.org/en/publication/international-humanitarian-law-and-challenges-contemporary-armed-conflicts-building. 1 July 2025.

Johnson, G. M. (2021). Algorithmic bias: On the implicit biases of social technology. *Synthese*, *198*, 9941–9961.

Johnson, G. M. (2023). Are Algorithms Value-Free? Feminist theoretical virtues in machine learning. *Journal of Moral Philosophy*, *21*(1–2), 27–61. https://doi.org/10.1163/17455243-20234372

Johnson, J. (2022). The AI commander problem: Ethical, political, and psychological dilemmas of human-machine interactions in AI-enabled Warfare. *Journal of Military Ethics*, *21*(3–4), 246–271. https://doi.org/10.1080/15027570.2023.2175887

Jones, N. A., Ross, H., Lynam, T., Perez, P., & Leitch, A. (2011). Mental models: An interdisciplinary synthesis of theory and methods. *Ecology and Society*, *16*(1), 1–13. https://www.ecologyandsociety.org/vol16/iss1/art46/

Kahneman, D. (2011). *Thinking fast and slow*. New York, NY, USA: Farrar, Straus & Giroux.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, *64*(6), 515–526. https://doi.org/10.1037/a0016755

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

Kares, F., König, C. J., Bergs, R., Protozel, C., & Langer, M. (2023). Trust in hybrid human-automated decision-support. *International Journal of Selection and Assessment*, *31*(1), 388–402. https://doi.org/10.1111/ijsa.12436

Kensinger, E. A., & Schacter, D. L. (2006). When the Red Sox shocked the Yankees: Comparing negative and positive memories. *Psychonomic Bulletin & Review*, *13*(5), 757–763. https://doi.org/10.3758/BF03193993

Klein, G. (2007). *The power of intuition: How to use your gut feelings to make better decisions at work*. New York, NY, USA: Doubleday Publishing Group.

Korteling, J. E., & Toet, A. (2021). Cognitive biases. In R. F. Kretschmann, (Ed.) *Encyclopedia of behavioral neuroscience* (2nd ed., Vol. 3, pp. 610–619). North Holland, Netherlands: Elsevier. https://doi.org/10.1016/B978-0-12-809324-5.24105-9.

Lee, S. J., Paas, L., & Ahn, H. S. (2024). The power of specific emotion analysis in predicting donations: A comparative empirical study between sentiment and specific emotion analysis in social media. *International Journal of Market Research*, *66*(5), 610–630. https://doi.org/10.1177/14707853241261248

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, *13*(5), 585–589. https://doi.org/10.1016/S0022-5371(74)80011-3

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press

Masalonis, A. J. (2003). Effects of training operators on situation-specific automation reliability. In *Proceedings of the 2003 IEEE International Conference on Systems, Man and Cybernetics*, *2*, 1595–1599. https://doi.org/10.1109/ICSMC.2003.1244115

Mueller, S., & Kuczynski, G. (2023). *Military decision-making process: Organizing and conducting planning*. USA: Center for Army Lessons Learned. https://api.army.mil/e2/c/downloads/2023/11/17/f7177a3c/23-07-594-military-decision-making-process-nov-23-public.pdf

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

Noble, S. U. (2018). *Algorithms of Oppression: How search engines reinforce racism*. New York: New York University Press.

Ntoutsi, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejdi, W., Vidal, M., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., Kompatsiaris,I., Kinder-Kurlanda, K., Wagner, C., Karimi, F., Fernandez, M., Alani, H., Berendt, B., Kruegel, T., … Staab, S. (2020). Bias in data-driven artificial intelligence systems – An introductory survey. *WIREs Data Mining and Knowledge Discovery*, *10*, e1356. https://doi.org/10.1002/widm.1356

Persson, E., Barrafrem, K., Meunier, A., & Tinghög, G. (2019). The effect of decision fatigue on surgeons' clinical decision making. *Health Economics*, *28*(10), 1194–1203. https://doi.org/10.1002/hec.3933

Phillips, J. K., Klein, G., & Sieck, W. R. (2004). Expertise in judgment and decision making: A case for training intuitive decision skills. In D. J. Koehler and N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 297–315). New Jersey, USA: Blackwell Publishing Ltd.

Pohl, R. F. (Ed.). (2004). *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory* (1st ed.). London, UK: Psychology Press. https://doi.org/10.4324/9780203720615

Prakasha, K. K., & Sumalatha, U. (2025). Privacy-preserving techniques in biometric systems: Approaches and challenges. *IEEE Access*, *13*, 32584–32616. https://doi.org/10.1109/ACCESS.2025.3541649

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 803–814. https://doi.org/10.1037/0278-7393.21.4.803

Shin, D., & Shin, E. Y. (2023). Data's impact on algorithmic bias. *Computer*, *56*(6), 90–94. https://doi.org/10.1109/MC.2023.3262909

Shin, M., Kim, J., van Opheusden, B., & Griffiths, T. L. (2023). Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences*, *120*(12), 1–6.e2214840120. https://doi.org/10.1073/pnas.2214840120

Sienknecht, M. (2024). The Prospect of Proxy Responsibility: Addressing Responsibility Gaps in Human-Machine Decision Making on the Resort to Force. *Australian Journal of International Affairs*, *78*(2), 191–199.

Simon, H. A. (1989). *Models of thought*. New Haven, Connecticut, USA: Yale University Press.

Sjøgren, S. (2022). What military commanders do and how they do it: Executive decision- making in the context of standardised planning processes and doctrine. *Scandinavian Journal of Military Studies*, *5*(1), 379–397. https://doi.org/10.31374/sjms.146

Skitka, L. J., Mosier, K., & Burdick, M. D. (2000). Accountability and automation bias. *International Journal of Human-Computer Studies*, *52*(4), 701–717. https://doi.org/10.1006/ijhc.1999.0349

Skitka, L. J., Mosier, K. L., & Burdick, M. (1999). Does automation bias decision-making? *International Journal of Human-Computer Studies*, *51*(5), 991–1006. https://doi.org/10.1006/ijhc.1999.0252

Sparrow, R. J., & Henschke, A. (2023). Minotaurs, not centaurs: The future of manned- unmanned teaming. *Parameters*, *53*(1), 115–130.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458. https://doi.org/10.1126/science.7455683

Vold, K. (2024). Human-AI cognitive teaming: Using AI to support state-level decision making on the resort to force. *Australian Journal of International Affairs*, *78*(2), 229–236. https://doi.org/10.1080/10357718.2024.2327383

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*(3), 273–281. https://doi.org/10.1080/14640746808400161

Wheeler, N., & Holmes, M. (2024). The role of artificial intelligence in nuclear crisis decision making. *Australian Journal of International Affairs*, *78*(2), 164–174. https://doi.org/10.1080/10357718.2024.2333814

Zerilli, J., Bhatt, U., & Weller, A. (2022). How transparency modulates trust in artificial intelligence. *Patterns*, *3*(4), 1–10. https://doi.org/10.1016/j.patter.2022.100455

Zerilli, J., Knott, A., Maclaurin, J., & Gavaghan, C. (2019). Algorithmic decision-making and the control problem. *Minds & Machines*, *29*(4), 555–578. https://doi.org/10.1007/s11023-019-09513-7

Cite this article: Vold, K. (2025). Augmenting military decision making with artificial intelligence. *Cambridge Forum on AI: Law and Governance*, *1*, e48, 1–13. https://doi.org/10.1017/cfl.2025.10022

https://doi.org/10.1017/cfl.2025.10022 Published online by Cambridge University Press