

Prediction and Judgment

Avi Goldfarb and
Jon R. Lindsay

Why Artificial Intelligence Increases the Importance of Humans in War

There is an emerging policy consensus that artificial intelligence (AI) will transform international politics. As stated in the 2021 report of the U.S. National Security Commission on AI, “The ability of a machine to perceive, evaluate, and act more quickly and accurately than a human represents a competitive advantage in any field—civilian or military. AI technologies will be a source of enormous power for the companies and countries that harness them.”¹ A lack of clarity over basic concepts, however, complicates an assessment of the security implications of AI. AI also has multiple meanings, ranging from big data, machine learning, robotics, and lethal drones, to a sweeping “fourth industrial revolu-

Avi Goldfarb is the Rotman Chair in Artificial Intelligence and Healthcare and Professor of Marketing at the University of Toronto, and a research associate at the National Bureau of Economic Research. Jon R. Lindsay is Associate Professor at the School of Cybersecurity and Privacy, with a secondary appointment at the Sam Nunn School of International Affairs, both at the Georgia Institute of Technology.

The authors are grateful for research assistance from Morgan MacInnes and constructive feedback from Andrea Gilli, Mauro Gilli, James Johnson, Ryder McKeown, Heather Roff, members of the Innovation Policy Lab at the Munk School of Global Affairs and Public Policy at the University of Toronto, and the anonymous reviewers. This project was supported by funding from the Social Sciences and Humanities Research Council of Canada (File number 435-2017-0041) and the Sloan Foundation. The authors presented a more limited version of this general argument in a previous report published by the Brookings Institution: Avi Goldfarb and Jon R. Lindsay, “Artificial Intelligence in War: Human Judgment as an Organizational Strength and a Strategic Liability” (Washington, D.C.: Brookings Institution Press, November 30, 2020), <https://www.brookings.edu/research/artificial-intelligence-in-war-human-judgment-as-an-organizational-strength-and-a-strategic-liability/>.

1. National Security Commission on Artificial Intelligence [NSCAI], *Final Report* (Washington, D.C.: NSAI, March 2021), p. 7, <https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf>. Russian President Vladimir Putin makes the same point with more flair: “Whoever becomes the leader in this sphere will become the ruler of the world,” quoted in Radina Gigova, “Who Putin Thinks Will Rule the World,” *CNN*, September 2, 2017, <https://www.cnn.com/2017/09/01/world/putin-artificial-intelligence-will-rule-world/index.html>; and Keith Dear, “Will Russia Rule the World through AI? Assessing Putin’s Rhetoric against Russia’s Reality,” *RUSI Journal*, Vol. 164, No. 5–6 (2019), pp. 36–60, <https://doi.org/10.1080/03071847.2019.1694227>. Many U.S. officials, such as the chief technology officer of the National Geospatial Intelligence Agency, draw a stark conclusion: “If the United States refuses to evolve, it risks giving China or some other adversary a technological edge that Washington won’t be able to overcome,” quoted in Anthony Vinci, “The Coming Revolution in Intelligence Affairs,” *Foreign Affairs*, August 31, 2020, <https://www.foreignaffairs.com/articles/north-america/2020-08-31/coming-revolution-intelligence-affairs>. On Chinese ambitions for an “intelligentized” military, see Elsa B. Kania, “Chinese Military Innovation in the AI Revolution,” *RUSI Journal*, Vol. 164, No. 5–6 (2019), pp. 26–34, <https://doi.org/10.1080/03071847.2019.1693803>.

International Security, Vol. 46, No. 3 (Winter 2021/22), pp. 7–50, https://doi.org/10.1162/isec_a_00425
© 2022 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

tion.”² Significant investment in AI and the imaginaries of science fiction only add to the confusion.

In this article we focus on machine learning, which is the key AI technology that receives attention in the press, in management, and in the economics literature.³ We leave aside debates about artificial general intelligence (AGI), or systems that match or exceed human intelligence.⁴ Machine learning, or “narrow AI,” by contrast, is already widely in use. Successful civil applications include navigation and route planning, image recognition and text translation, and targeted advertising. Michael Horowitz describes AI as “the ultimate enabler” for automating decision-making tasks in everything from public administration and commercial business to strategic intelligence and military combat.⁵ In 2018, the Department of Defense observed that “AI is poised to transform every industry, and is expected to impact every corner of the Department, spanning operations, training, sustainment, force protection, recruiting, healthcare, and many others.”⁶ We would be surprised, however, if AI transformed all these activities to the same degree for all actors who use it.

One of the key insights from the literature on the economics of technology is that the complements to a new technology determine its impact.⁷ AI, from this perspective, is not a simple substitute for human decision-making. Rapid advances in machine learning have improved statistical prediction, but predic-

2. On the potential impacts of these technologies, see the recent special issue edited by Michael Raska et al., “Introduction,” *Journal of Strategic Studies*, Vol. 44, No. 4 (2021), pp. 451–455, <https://doi.org/10.1080/01402390.2021.1917877>.

3. See, for example, Ajay Agrawal, Joshua Gans, and Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence* (Cambridge, Mass.: Harvard Business Review Press, 2018), p. 24; and Jason Furman and Robert Seamans, “AI and the Economy,” in Josh Lerner and Scott Stern, eds., *Innovation Policy and the Economy*, Vol. 19 (Chicago: University of Chicago Press, 2018), pp. 161–191.

4. Generally, see Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014). We comment briefly on AGI in the conclusion.

5. Michael C. Horowitz, “Artificial Intelligence, International Competition, and the Balance of Power,” *Texas National Security Review*, Vol. 1, No. 3 (May 2018), p. 41, <https://doi.org/10.15781/T2639KP49>.

6. *Summary of the 2018 Department of Defense Artificial Intelligence Strategy: Harnessing AI to Advance Our Security and Prosperity* (Washington, D.C.: U.S. Department of Defense, February 12, 2019), p. 5, <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF>. On potential military applications, see NSCAI, *Final Report*; and Daniel S. Hoadley and Kelley M. Saylor, *Artificial Intelligence and National Security*, CRS Report R45178 (Washington, D.C.: Congressional Research Service, November 21, 2019), <https://crsreports.congress.gov/product/pdf/R/R45178/7>.

7. See, for example, Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt, “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *Quarterly Journal of Economics*, Vol. 117, No. 1 (February 2002), pp. 339–376, <https://doi.org/10.1162/003355302753399526>; and Shane Greenstein and Timothy F. Bresnahan, “Technical Progress and Co-invention in Computing and in the Uses of Computers,” *Brookings Papers on Economic Activity: Microeconomics* (Washington, D.C.: Brookings Institution Press, 1996), pp. 1–83.

tion is only one aspect of decision-making. Two other important elements of decision-making—data and judgment—represent the complements to prediction. Just as cheaper bread expands the market for butter, advances in AI that reduce the costs of prediction are making its complements more valuable. AI prediction models require data, and accurate prediction requires more and better data. Quality data provide plentiful and relevant information without systemic bias. Data-driven machine prediction can efficiently fill in information needed to optimize a given utility function, but the specification of the utility function ultimately relies on human judgment about what exactly should be maximized or minimized. Judgment determines what kinds of patterns and outcomes are meaningful and what is at stake, for whom, and in which contexts. Clear judgments are well specified in advance and agreed upon by relevant stakeholders. When quality data are available and an organization can articulate clear judgments, then AI can improve decision-making.

We argue that if AI makes prediction cheaper for military organizations, then data and judgment will become both more valuable and more contested. This argument has two important strategic implications. First, the conditions that have made AI successful in the commercial world—quality data and clear judgment—may not be present, or present to the same degree, for all military tasks. In military terms, judgment encompasses command intentions, rules of engagement, administrative management, and moral leadership. These functions cannot be automated with narrow AI technology. Increasing reliance on AI, therefore, will make human beings even more vital for military power, not less. Second, the importance of data and judgment creates incentives for strategic competitors to improve, protect, and interfere with information systems and command institutions. As a result, conflicts over information will become more salient, and organizational coordination will become more complex. In contrast with assumptions about rapid robot wars and decisive shifts in military advantage, we expect AI-enabled conflict to be characterized by environmental uncertainty, organizational friction, and political controversy. The contestation of AI complements, therefore, is likely to unfold differently than the imagined wars of AI substitutes.⁸

Many hopes and fears about AI recapitulate earlier ideas about the in-

8. See, for example, Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W.W. Norton, 2018); Michael C. Horowitz, "When Speed Kills: Lethal Autonomous Weapon Systems, Deterrence and Stability," *Journal of Strategic Studies*, Vol. 42, No. 6 (2019), pp. 764–788, <https://doi.org/10.1080/01402390.2019.1621174>; James Johnson, "Artificial Intelligence and Future Warfare: Implications for International Security," *Defense & Security Analysis*, Vol. 35, No. 2 (2019), pp. 147–169, <https://doi.org/10.1080/14751798.2019.1600800>; and Kenneth Payne, *I, Warbot: The Dawn of Artificially Intelligent Conflict* (New York: Oxford University Press, 2021).

formation technology revolution in military affairs (RMA) and cyberwarfare.⁹ Familiar tropes abound regarding the transformative effects of commercial innovation, the speed and danger of networked computation, the dominance of offense over defense, and the advantages of a rising China over a vulnerable United States. But skeptics have systematically challenged both the logic and empirical basis of these assumptions about the RMA¹⁰ and cyberwar.¹¹ Superficially plausible arguments about information technology tend to ignore important organizational and strategic factors that shape the adoption and use of digital systems. As in the economics literature, an overarching theme in scholarship on military innovation is that technology is not a simple substitute for military power.¹² Technological capabilities depend on complementary institutions, skills, and doctrines. Furthermore, implementation is usually marked by friction, unintended consequences, and disappointed expectations. The RMA and cyber debates thus offer a cautionary tale for claims about AI. It is reasonable to expect organizational and strategic context to condition the performance of automated systems, as with any other information technology.¹³

AI may seem different, nevertheless, because human agency is at stake. Recent scholarship raises a host of questions about the prospect of auto-

9. See, for example, John Arquilla and David Ronfeldt, "Cyberwar Is Coming!," *Comparative Strategy*, Vol. 12, No. 2 (Spring 1993), pp. 141–165; Arthur K. Cebrowski and John H. Garstka, "Network-Centric Warfare: Its Origin and Future," *Proceedings*, U.S. Naval Institute, January, 1998, pp. 28–35; William A. Owens and Edward Offley, *Lifting the Fog of War* (New York: Farrar, Straus and Giroux, 2000); and Richard A. Clarke and Robert K. Knake, *Cyber War: The Next Threat to National Security and What to Do about It* (New York: Ecco, 2010).

10. See, for example, Eliot A. Cohen, "Change and Transformation in Military Affairs," *Journal of Strategic Studies*, Vol. 27, No. 3 (2004), pp. 395–407, <https://doi.org/10.1080/1362369042000283958>; and Keith L. Shimko, *The Iraq Wars and America's Military Revolution* (New York: Cambridge University Press, 2010).

11. See, for example, Erik Gartzke, "The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth," *International Security*, Vol. 38, No. 2 (Fall 2013), pp. 41–73, https://doi.org/10.1162/ISEC_a_00136; Jon R. Lindsay, "The Impact of China on Cybersecurity: Fiction and Friction," *International Security*, Vol. 39, No. 3 (Winter 2014/15), pp. 7–47, https://doi.org/10.1162/ISEC_a_00189; Brandon Valeriano and Ryan C. Maness, *Cyber War versus Cyber Realities: Cyber Conflict in the International System* (New York: Oxford University Press, 2015); and Rebecca Slayton, "What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment," *International Security*, Vol. 41, No. 3 (Winter 2016/17), pp. 72–109, https://doi.org/10.1162/ISEC_a_00267.

12. Reviews include Adam Grissom, "The Future of Military Innovation Studies," *Journal of Strategic Studies*, Vol. 29, No. 5 (2006), pp. 905–934, <https://doi.org/10.1080/01402390600901067>; and Michael C. Horowitz, "Do Emerging Military Technologies Matter for International Politics?" *Annual Review of Political Science*, Vol. 23 (May 2020), pp. 385–400, <https://doi.org/10.1146/annurev-polisci-050718-032725>. See, in particular, Stephen Biddle, *Military Power: Explaining Victory and Defeat in Modern Battle* (Princeton, N.J.: Princeton University Press, 2004); Michael C. Horowitz, *The Diffusion of Military Power: Causes and Consequences for International Politics* (Princeton, N.J.: Princeton University Press, 2010); and Andrea Gilli and Mauro Gilli, "Why China Has Not Caught Up Yet: Military-Technological Superiority and the Limits of Imitation, Reverse Engineering, and Cyber Espionage," *International Security*, Vol. 43, No. 3 (Winter 2018/19), pp. 141–189, https://doi.org/10.1162/isec_a_00337.

13. On general patterns of information practice in war, see Jon R. Lindsay, *Information Technology and Military Power* (Ithaca, N.Y.: Cornell University Press, 2020).

mated decision-making. How will war “at machine speed” transform the offense-defense balance?¹⁴ Will AI undermine deterrence and strategic stability,¹⁵ or violate human rights?¹⁶ How will nations and coalitions maintain control of automated warriors?¹⁷ Does AI shift the balance of power from incumbents to challengers or from democracies to autocracies?¹⁸ These questions focus on the substitutes for AI because they address the political, operational, and moral consequences of replacing people, machines, and processes with automated systems. The literature on military AI has focused less on the complements of AI, namely the organizational infrastructure, human skills, doctrinal concepts, and command relationships that are needed to harness the advantages and mitigate the risks of automated decision-making.¹⁹

14. Kenneth Payne, “Artificial Intelligence: A Revolution in Strategic Affairs?” *Survival*, Vol. 60, No. 5 (2018), pp. 7–32, <https://doi.org/10.1080/00396338.2018.1518374>; Paul Scharre, “How Swarming Will Change Warfare,” *Bulletin of the Atomic Scientists*, Vol. 74, No. 6 (2018), pp. 385–389, <https://doi.org/10.1080/00963402.2018.1533209>; Ben Garfinkel and Allan Dafoe, “How Does the Offense-Defense Balance Scale?” *Journal of Strategic Studies*, Vol. 42, No. 6 (2019), pp. 736–763, <https://doi.org/10.1080/01402390.2019.1631810>; and John R. Allen, Frederick Ben Hodges, and Julian Lindley-French, “Hyperwar: Europe’s Digital and Nuclear Flanks,” in Allen, Hodges, and Lindley-French, *Future War and the Defence of Europe* (New York: Oxford University Press, 2021), pp. 216–245.

15. Jürgen Altmann and Frank Sauer, “Autonomous Weapon Systems and Strategic Stability,” *Survival*, Vol. 59, No. 5 (2017), pp. 117–142, <https://doi.org/10.1080/00396338.2017.1375263>; Horowitz, “When Speed Kills”; Mark Fitzpatrick, “Artificial Intelligence and Nuclear Command and Control,” *Survival*, Vol. 61, No. 3 (2019), pp. 81–92, <https://doi.org/10.1080/00396338.2019.1614782>; Erik Gartzke, “Blood and Robots: How Remotely Piloted Vehicles and Related Technologies Affect the Politics of Violence,” *Journal of Strategic Studies*, published online October 3, 2019, <https://doi.org/10.1080/01402390.2019.1643329>; and James Johnson, “Delegating Strategic Decision-Making to Machines: Dr. Strangelove Redux?” *Journal of Strategic Studies*, published online April 30, 2020, <https://doi.org/10.1080/01402390.2020.1759038>.

16. Ian G.R. Shaw, “Robot Wars: US Empire and Geopolitics in the Robotic Age,” *Security Dialogue*, Vol. 48, No. 5 (2017), pp. 451–470, <https://doi.org/10.1177/0967010617713157>; and Lucy Suchman, “Algorithmic Warfare and the Reinvention of Accuracy,” *Critical Studies on Security*, Vol. 8, No. 2 (2020), pp. 175–187, <https://doi.org/10.1080/21624887.2020.1760587>.

17. Heather M. Roff, “The Strategic Robot Problem: Lethal Autonomous Weapons in War,” *Journal of Military Ethics*, Vol. 13, No. 3 (2014), pp. 211–227, <https://doi.org/10.1080/15027570.2014.975010>; Heather M. Roff and David Danks, “‘Trust but Verify’: The Difficulty of Trusting Autonomous Weapons Systems,” *Journal of Military Ethics*, Vol. 17, No. 1 (2018), pp. 2–20, <https://doi.org/10.1080/15027570.2018.1481907>; Risa Brooks, “Technology and Future War Will Test U.S. Civil-Military Relations,” *War on the Rocks*, November 26, 2018, <https://warontherocks.com/2018/11/technology-and-future-war-will-test-u-s-civil-military-relations/>; and Erik Lin-Greenberg, “Allies and Artificial Intelligence: Obstacles to Operations and Decision-Making,” *Texas National Security Review*, Vol. 3, No. 2 (Spring 2020), pp. 56–76, <https://dx.doi.org/10.26153/tsw/8866>.

18. Horowitz, “Artificial Intelligence, International Competition, and the Balance of Power”; Ben Buchanan, “The U.S. Has AI Competition All Wrong,” *Foreign Affairs*, August 7, 2020, <https://www.foreignaffairs.com/articles/united-states/2020-08-07/us-has-ai-competition-all-wrong>; and Michael Raska, “The Sixth RMA Wave: Disruption in Military Affairs?” *Journal of Strategic Studies*, Vol. 44, No. 4 (2021), pp. 456–479, <https://doi.org/10.1080/01402390.2020.1848818>.

19. A notable exception is Horowitz, “Artificial Intelligence, International Competition, and the Balance of Power.” We agree with Horowitz that organizational complements determine AI diffusion, but we further argue that complements also shape AI employment, which leads us to different expectations about future war.

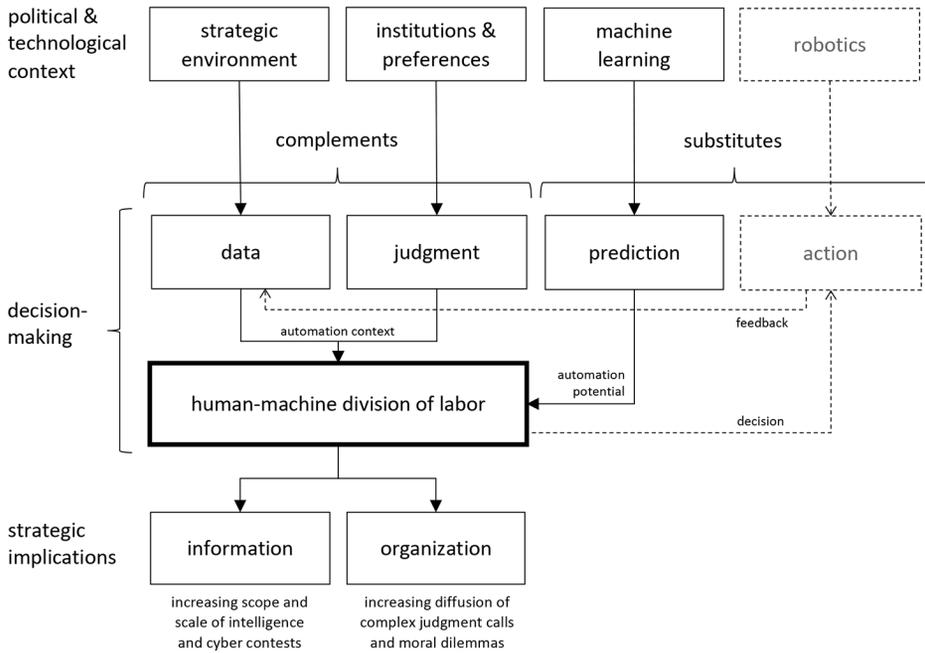
In this article, we challenge the assumptions behind AI substitution and explore the implications of AI complements. An army of lethal autonomous weapon systems may well be destabilizing, and such an army may well be attractive to democracies and autocracies alike. The idea that machines will replace warriors, however, represents a misunderstanding about what warriors actually do. We suggest that it is premature to forecast radical strategic consequences without first clarifying the problem that AI is supposed to solve. We provide a framework that explains how the complements of AI (i.e., data and judgment) affect decision-making. In general, automation is advantageous when quality data can be combined with clear judgments. But the consummate military tasks of command, fire, and maneuver are fraught with uncertainty and confusion. In contrast, more institutionalized tasks in administration and logistics tend to have copious data and clear goals, which are conducive to automation. We argue that militaries risk facing bad or tragic outcomes if they conflate these conditions by prematurely providing autonomous systems with clear objectives in uncertain circumstances. Conversely, for intelligence and operational tasks that have quality data but difficult judgments, teams of humans and machines can distribute the cognitive load of decision-making. We expect many if not most military AI tasks to fall into the latter category, which we describe as human-machine teaming. The net result, we argue, is that data and judgment will become increasingly valuable and contested, and thus AI-enabled warfare will tend to become more protracted and confusing.

We develop this argument in five parts. First, we provide an overview of our analytical framework, which distinguishes the universal process of decision-making from its variable political and technological context. This framework explains how data and judgment affect the human-machine division of labor in decision-making. Second, we describe how strategic and institutional conditions, which differ in business and military affairs, shape the quality of data and the difficulty of judgment. We then combine these factors into four different categories of AI performance of decision-making tasks, which we illustrate with commercial and military examples. The penultimate section discusses the strategic implications of data and judgment becoming more valuable. We conclude with a summary of the argument and further implications.

The Political and Technological Context of Decision-Making

Business and military organizations are similar in many ways, but they operate in very different circumstances. In figure 1, we analytically distinguish the AI-relevant similarities and differences by embedding an economic model of

Figure 1. The Strategic Context of Decision-Making in Military Organizations



decision-making into an international relations framework.²⁰ Decision-making both shapes and is shaped by the political and technological context. The strategic environment and organizational institutions affect the quality of data and judgment, respectively. At the same time, innovation in machine learning—largely driven by the civilian sector—lowers the cost of prediction.

20. Elements depicted in dashed lines in figure 1 are important for the overall story, but we will not discuss them in detail in this article. We include them to depict the full decision-making process—data, judgment, prediction, action—and to distinguish machine learning from other types of automation technology. We thus discuss robotics or drones that automate military action only to the extent that machine learning provides a decision input for them. Similar considerations about complementarity apply to drones as well; see Andrea Gilli and Mauro Gilli, “The Diffusion of Drone Warfare? Industrial, Organizational, and Infrastructural Constraints,” *Security Studies*, Vol. 25, No. 1 (2016), pp. 50–84, <https://doi.org/10.1080/09636412.2016.1134189>. For the sake of parsimony, we also omit intelligence, surveillance, and reconnaissance (ISR) technologies that affect data, as well as information and communication technologies (ICTs) that help coordinate anything whatsoever. Again, the logic of complementarity applies more generally to ICTs; see Lindsay, *Information Technology and Military Power*. While our focus in this article is on theory building rather than testing, the same framework in figure 1 could be used to compare cases (e.g., in business, the military, or cross-nationally) that leverage similar AI technology but in different contexts.

POLITICAL CONTEXT: ENVIRONMENT AND INSTITUTIONS

We adopt standard international relations distinctions between the international system and domestic institutions.²¹ The “strategic environment” in figure 1 refers to the external problems confronting a military organization. To alter or preserve facts on the ground, through conquest or denial, a military needs information about many things, such as the international balance of power, diplomatic alignments and coalitions, geographical terrain and weather, the enemy’s operational capabilities and disposition, and interactions with civil society. These external matters constitute threats, targets, opportunities, resources, and constraints for military operations. A military also needs information about internal matters, such as the capabilities and activities of friendly forces, but these are a means to an end.²² The strategic environment is ultimately what military data are about, and the structure and dynamics of the environment affect the quality of the data.

“Institutions and preferences” in figure 1 refer to the ways in which a military organization solves its strategic problems. This general category encompasses bureaucratic structures and processes, interservice and coalition politics, civil-military relations, interactions with the defense industry, and other domestic politics. Any of these factors might influence the goals and values of a military organization or the way it interprets a given situation. Organizational institutions embody preferences, whatever their source, which in turn

21. Our modest goal here is to emphasize that strategic and organizational context shapes the performance of AI technology in military decision-making tasks. In this article we do not take a position on which of these contextual factors will be more influential in any given situation. We also omit heterogeneity within and dynamic interactions across these factors. Future research could disaggregate these factors to explore more specific hypotheses about AI and military power. Lindsay, *Information Technology and Military Power*, pp. 32–70, discusses a similar framework in more detail but with different nomenclature. The general analytic distinction between environment, organizations, and technology is also employed by Barry R. Posen, *The Sources of Military Doctrine: France, Britain, and Germany between the World Wars* (Ithaca, N.Y.: Cornell University Press, 1984). On the interaction between system- and unit-level factors in realist theory, see Gideon Rose, “Neoclassical Realism and Theories of Foreign Policy,” *World Politics*, Vol. 51, No. 1 (October 1998), pp. 144–172, <https://doi.org/10.1017/S0043887100007814>; and Kevin Narizny, “On Systemic Paradigms and Domestic Politics: A Critique of the Newest Realism,” *International Security*, Vol. 42, No. 2 (Fall 2017), pp. 155–190, https://doi.org/10.1162/ISEC_a_00296.

22. From a rationalist perspective, whereby preferences are exogenously specified, internal processes are instrumental to the fundamental task of knowing and influencing the world. In natural systems, of course, internal processes may become infused with value and endogenously shape organizational behavior, as discussed by Philip Selznick, *TVA and the Grass Roots: A Study in the Sociology of Formal Organization* (Berkeley: University of California Press, 1949); and Herbert Kaufman, *The Forest Ranger: A Study in Administrative Behavior* (Baltimore, Md.: John Hopkins University Press, 1960). Our framework allows for both possibilities by interacting environmental data and organizational preferences, whatever their source. Furthermore, at the level of analysis of any given decision task, the “environment” can be analyzed as including “institutions” too. We omit these complexifying relationships because they only reinforce our general point about the importance of context.

affect the quality of judgment.²³ Institutional structures and processes may produce coordination problems, political controversies, or interpretive difficulties that make it hard for a military organization to figure out what matters and why.

Furthermore, as discussed below, we expect that the adoption of AI for some military decision tasks will (endogenously) affect the strategic environment and military institutions over time. As data and judgment become more valuable, strategic competitors will have incentives to improve and contest them. We thus expect conflicts over information to become more salient while organizational coordination will become more complex.

TECHNOLOGICAL CONTEXT: MACHINE LEARNING AS PREDICTION

The resurgence of interest in AI since the turn of the millennium has been driven by rapid advances in a subfield called machine learning. Machine learning techniques represent a different approach compared to “Good Old-Fashioned AI” (GOFAI).²⁴ GOFAI emphasizes deductive theorem proving and search optimization. Machine learning, by contrast, is a form of statistical prediction, which is the process of using existing data to inductively generate missing information.²⁵ While the term prediction often implies forecasting the future, pattern recognition and object classification are also forms of prediction because they fill in information about situations encountered for the first time. Machines can automate many prediction tasks that humans perform today (e.g., image recognition, navigation, and forecasting), and they can also increase the number, accuracy, complexity, and speed of predictions. This has the potential to alter human workflows. While machines may not make decisions, they can alter who makes what decisions and when. As machine learning lowers the cost of prediction, organizations are also innovating ways to improve data and judgments so that they can make better decisions.

Prediction usually involves generalizing from a set of training data to clas-

23. Internal processes and external situations may affect the goals and values of an organization. Our framework is agnostic regarding the ultimate source of preferences. Whatever their source, preferences become embodied in institutions that proximally transmit goals and values to decision-makers. On the general debate between structuralists and institutionalists about political and doctrinal preferences, see Posen, *The Sources of Military Doctrine*; Alexander E. Wendt, “The Agent-Structure Problem in International Relations Theory,” *International Organization*, Vol. 41, No. 3 (Summer 1987), pp. 335–370, <http://www.jstor.org/stable/2706749>; Stephen van Evera, “Hypotheses on Nationalism and War,” *International Security*, Vol. 18, No. 4 (Spring 1994), pp. 5–39, <https://doi.org/10.2307/2539176>; and Jeffrey W. Legro and Andrew Moravcsik, “Is Anybody Still a Realist?” *International Security*, Vol. 24, No. 2 (Fall 1999), pp. 5–55, <https://doi.org/10.1162/016228899560130>.

24. Terrence J. Sejnowski, *The Deep Learning Revolution* (Cambridge: Massachusetts Institute of Technology Press, 2018).

25. Agrawal, Gans, and Goldfarb, *Prediction Machines*, p. 24.

sify or synthesize new data. It became clear in the early 2000s that improvements in computing, memory, and bandwidth could make machine learning commercially viable. Firms like Google, Amazon, and Facebook have successfully targeted their advertising and digital services by coupling “big data” that they harvest from consumer behavior with automated prediction techniques. These same developments have also enabled espionage and surveillance at an unprecedented scale.²⁶

From an economic perspective, modern AI is best understood as a better, faster, and cheaper form of statistical prediction. The overall effect on decision-making, however, is indeterminate. This implies that organizations, military and otherwise, will be able to perform more prediction in the future than they do today, but not necessarily that their performance will improve in all cases.

THE DECISION-MAKING PROCESS

Economic decision theory emerged alongside the intellectual tradition of cybernetics.²⁷ As Herbert Simon observed over sixty years ago, “A real-life decision involves some goals or values, some facts about the environment, and some inferences drawn from the values and facts.”²⁸ We describe these elements as judgment, data, and prediction. Together they produce actions that shape economic or political outcomes. Feedback from actions produces more data, which can be used for more predictions and decisions, or to reinterpret judgment. The so-called OODA loop in military doctrine captures the same ideas.²⁹ Decision cycles govern all kinds of decision tasks, from the trivial (picking up a pencil) to the profound (mobilizing for war). The abstract

26. David V. Goe, Michael S. Goodman, and Tim Stevens, “Intelligence in the Cyber Era: Evolution or Revolution?” *Political Science Quarterly*, Vol. 135, No. 2 (Summer 2020), pp. 191–224, <https://doi.org/10.1002/polq.13031>.

27. The main ideas from the economics literature on decision-making are summarized in Itzhak Gilboa, *Making Better Decisions: Decision Theory in Practice* (Oxford: Wiley-Blackwell, 2011). See also John D. Steinbruner, *The Cybernetic Theory of Decision: New Dimensions of Political Analysis* (Princeton, N.J.: Princeton University Press, 1974). On the intellectual impact of cybernetics generally, see Ronald R. Kline, *The Cybernetics Moment: Or Why We Call Our Age the Information Age* (Baltimore, Md.: Johns Hopkins University Press, 2015). Classic applications of cybernetic decision theory include Karl W. Deutsch, *The Nerves of Government: Models of Political Communication and Control* (New York: Free Press, 1963); and James R. Beniger, *The Control Revolution: Technological and Economic Origins of the Information Society* (Cambridge, Mass.: Harvard University Press, 1989).

28. Herbert A. Simon, “Theories of Decision-Making in Economics and Behavioral Science,” *American Economic Review*, Vol. 49, No. 3 (June 1959), p. 273, <https://www.jstor.org/stable/1809901>.

29. OODA stands for the “observe, orient, decide, and act” phases of the decision cycle. Note that “orient” and “decide” map to prediction and judgment, respectively. These phases may occur sequentially or in parallel in any given implementation. On the influence of John Boyd’s cybernetic OODA loop in military thought see James Hasik, “Beyond the Briefing: Theoretical and Practical Problems in the Works and Legacy of John Boyd,” *Contemporary Security Policy*, Vol. 34, No. 3 (2013), pp. 583–599, <https://doi.org/10.1080/13523260.2013.839257>.

decision model is agnostic about implementation, which means that the logic of decision might be implemented with organic, organizational, or technological components.

Yet, implementation is precisely what is at stake with AI. Figure 1 thus illustrates how data, judgment, prediction, and action affect the human-machine division of labor in decision-making. We highlight the human-machine division of labor in decision-making because the task-specific implementation of AI has significant consequences. The universality of cybernetic decision-making explains why many consider AI to be a general-purpose technology, like electricity or the internal combustion engine.³⁰ AI can indeed improve prediction, which is a vital input for any sort of decision-making. But AI is not the only input. Organizations also rely on data and judgment to make decisions in task-specific circumstances. Put simply, AI is a general-purpose technology that performs differently in specific contexts.

THE HUMAN-MACHINE DIVISION OF LABOR

When analyzing or deploying AI, it is necessary to consider particular tasks that serve particular goals. Machine learning is not AGI. Our unit of analysis, therefore, is the decision task, that is, the ensemble of data, predictions, judgments, and actions that produce a specific organizational outcome. Most organizations perform many different and interrelated tasks, such as strategy, management, human resources, marketing, network administration, manufacturing, operations, security, and logistics. Military analogs of these tasks include command, administration, training, intelligence, communication, fire, maneuver, protection, and sustainment. Within and across these categories are myriad different tactics, techniques, and procedures. Any of these tasks, at whatever scope or scale, can directly or indirectly support an organization's overall mission, which may or may not be well defined. Indeed, a task itself may be poorly defined, in part because task decomposition is a problem for managerial judgment.

AI performance in any given task is a function of the quality of data and the difficulty of judgment. These two complements provide essential context for automated decision-making. Data are high quality if relevant information is abundantly available and not systematically biased. Judgment is well defined if goals can be clearly specified in advance and stakeholders agree on them.

30. See, for example, Iain M. Cockburn, Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda* (Chicago: University of Chicago Press, 2019), pp. 115–146.

Table 1. The Implications of Data and Judgment for Automation in Decision-Making

		Data	
		high-quality	low-quality
Judgment	clear	Fully automated decision-making is more efficient.	Full automation is possible but risky.
	difficult	Automated predictions can inform human decisions.	Automated decision-making is not feasible.

The degree to which data quality is high or low, and judgment is clear or difficult, determines the comparative advantages of humans and machines in decision-making. In contrast, substitutes determine only the technical potential for automation (i.e., by reducing the costs of prediction or action). Table 1 summarizes the effects of data and judgment on AI performance. The implications of AI for future war are necessarily speculative, which makes it even more important to theorize from a solid deductive foundation.

AI Complements in Business and War

Numerous empirical studies of civilian AI systems have validated the basic finding that AI depends on quality data and clear judgment.³¹ To the extent that decision-making is indeed a universal activity, it is reasonable to expect models of it to apply to militaries. But we caution against generalizing from the business world to military affairs (and vice versa). Commercial and military organizations perform dissimilar tasks in different contexts. Militaries are only infrequently “in business” because wars are rare events.³² Objectives such as “victory” or “security” are harder to define than “shareholder value” or “profit.” Combatants attempt to physically destroy their competitors, and the consequences of failure are potentially existential.

One underappreciated reason why AI has been applied successfully in many commercial situations is because the enabling conditions of quality data and clear judgment are often present. Peaceful commerce generally takes place in institutionalized circumstances. Laws, property rights, contract enforcement mechanisms, diversified markets, common expectations, and shared behavioral norms all benefit buyers and sellers. These institutional features make

31. Detailed arguments and evidence are presented in Agrawal, Gans, and Goldfarb, *Prediction Machines*.

32. Gary King and Langche Zeng, “Explaining Rare Events in International Relations,” *International Organization*, Vol. 55, No. 3 (2001), pp. 693–715, <https://doi.org/10.1162/00208180152507597>.

transactions more consistent and efficient.³³ Consistency, in turn, provides the essential scaffolding for full automation. We expect AI to be more successful in more institutionalized circumstances and for more structured tasks.

War, by contrast, occurs in a more anarchic environment. In the international system, according to the intellectual tradition of realism, there are no legitimate overarching institutions to adjudicate disputes, enforce international agreements, or constrain behavior.³⁴ Actors must be prepared to defend themselves or ally with others for protection. Allies and adversaries alike have incentives to misrepresent their capabilities and interests, and for the same reasons to suspect deception by others.³⁵ Militarized crises and conflicts abound in secrecy and uncertainty. War aims are controversial, almost by definition, and they mobilize the passions of the nation, for better or worse.

We expect the absence of constraining institutions in war to undermine the AI-enabling conditions of quality data and clear judgment. One exception that proves the rule is that a military bureaucracy may be able to provide scaffolding for some military tasks. Robust organizational institutions, in other words, might substitute for weak international institutions. Yet, there are limits to what organizations can accomplish in the inherently uncertain and contested environment of war. The specific context of data and judgment will determine the viability of automation for any given task.

DATA AND THE STRATEGIC ENVIRONMENT

Commercial AI systems often need thousands, if not millions or billions, of examples to generate high-quality predictions. As deep-learning pioneer Geoffrey Hinton puts it, “Take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to

33. See, for example, R.H. Coase, “The Problem of Social Cost,” *Journal of Law and Economics*, Vol. 3 (October 1960), pp. 1–44, <https://doi.org/10.1086/466560>; and Oliver E. Williamson, “The Economics of Organization: The Transaction Cost Approach,” *American Journal of Sociology*, Vol. 87, No. 3 (November 1981), pp. 548–577, <https://www.jstor.org/stable/2778934>. These same features are associated with liberal perspectives on international politics, such as Robert O. Keohane, “The Demand for International Regimes,” *International Organization*, Vol. 36, No. 2 (Spring 1982), pp. 325–355, <https://www.jstor.org/stable/2706525>; John R. Oneal and Bruce M. Russett, “The Classical Liberals Were Right: Democracy, Interdependence, and Conflict, 1950–1985,” *International Studies Quarterly*, Vol. 41, No. 2 (June 1997), pp. 267–293, <https://doi.org/10.1111/1468-2478.00042>; and G. John Ikenberry, *After Victory: Institutions, Strategic Restraint, and the Rebuilding of Order after Major Wars* (Princeton, N.J.: Princeton University Press, 2001). A suggestive hypothesis that we do not develop in depth here is that the same conditions that are conducive to liberal institutions should also be conducive to AI performance.

34. These are standard assumptions in realist theory. For example, Hans J. Morgenthau, *Politics among Nations: The Struggle for Power and Peace* (New York: Alfred A. Knopf, 1960); and Kenneth N. Waltz, *Theory of International Politics* (Reading, Mass.: Addison-Wesley, 1979).

35. James D. Fearon, “Rationalist Explanations for War,” *International Organization*, Vol. 49, No. 3 (Summer 1995), pp. 379–414, <https://www.jstor.org/stable/2706903>.

make it work better than the existing techniques.”³⁶ The need for a large *quantity* of data, as well as detailed metadata that labels and describes its content, is well understood. But the need for *quality* data is less appreciated. Two factors can undermine the relevancy of data. First, the data may be biased toward one group of people or situation.³⁷ Second, data on the particular situation being predicted may not exist. This latter situation happens surprisingly often, because predictions are especially useful when they provide insight into what will happen if an organization changes its behavior. If the organization has never behaved in a certain way, then relevant data will not exist, and the related statistical prediction will fail.³⁸

A competitor or adversary can exacerbate both problems of data relevancy by manipulating data to create bias or interdicting the supply of data. If an adversary finds a way to access and corrupt the data used to train AI systems, then predictions become less reliable.³⁹ More generally, if AI becomes good at optimizing the solution for any given problem, then an intelligent enemy has incentives to change the problem. In the parlance of AI, the enemy will “go beyond the training set” by creating a situation for which there is no prior example in data used for machine learning. An adversary could innovate new tactics that are hard for AI systems to detect or pursue aims that AI systems do not anticipate.

The sources of uncertainty in war are legion (e.g., poor weather, unfamiliar terrain, bad or missing intelligence, misperception, enemy deception). As Carl von Clausewitz famously states, “War is the realm of uncertainty; three quarters of the factors on which action in war is based are wrapped in a fog of greater or lesser uncertainty.”⁴⁰ Clausewitz also uses the mechanical metaphor

36. Geoffrey Hinton, “On Radiology,” presented at the Machine Learning and Market for Intelligence Conference, Creative Destruction Lab, University of Toronto, October 26, 2016, YouTube video, 1:24, <https://youtu.be/2HMPRXstSvQ>.

37. Bo Cowgill and Catherine E. Tucker, “Algorithmic Fairness and Economics,” Columbia Business School Research Paper, SSRN (February 14, 2020), <https://dx.doi.org/10.2139/ssrn.3361280>; and Jon Kleinberg et al., “Discrimination in the Age of Algorithms,” *Journal of Legal Analysis*, Vol. 10 (2018), pp. 113–174, <https://doi.org/10.1093/jla/laz001>.

38. In such cases, causal inference requires different tools because the counterfactual situation is never observed. There is a rich technical literature on these ideas, rooted in the Rubin causal model. Widely used textbooks are Joshua D. Angrist and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist’s Companion* (Princeton, N.J.: Princeton University Press, 2009); and Guido W. Imbens and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction* (New York: Cambridge University Press, 2015).

39. Battista Biggio and Fabio Roli, “Wild Patterns: Ten Years after the Rise of Adversarial Machine Learning,” *Pattern Recognition*, Vol. 84 (December 2018), pp. 317–331, <https://doi.org/10.1016/j.patcog.2018.07.023>; and Heather M. Roff, “AI Deception: When Your Artificial Intelligence Learns to Lie,” *IEEE Spectrum*, February 24, 2020, <https://spectrum.ieee.org/automaton/artificial-intelligence/embedded-ai/ai-deception-when-your-ai-learns-to-lie>.

40. Carl von Clausewitz, *On War*, ed. and trans. Michael Eliot Howard and Peter Paret (Princeton, N.J.: Princeton University Press, 1989), p. 101.

of “friction” to describe organizational breakdowns. “Friction” in information technology can also create “fog” because the same systems adopted to improve certainty become new sources of uncertainty. Military personnel in networked organizations struggle to connect systems, negotiate data access, customize software, and protect information security. Computer glitches and configuration problems tend to “accumulate and end by producing a kind of friction that is inconceivable unless one has experienced war.”⁴¹ Empirical studies of wartime practice reveal a surprising amount of creative hacking to repair and reshape technologies to deal with the tremendous friction of warfare in the information age.⁴² Yet, informal adaptation of data processing systems can also create interoperability, accountability, and security problems. Managerial intervention to control these risks creates even more friction. We suggest that AI systems designed to “lift the fog of war” could just as easily “shift the fog” right back into the organization.⁴³

Although war is rife with uncertainties that can distort or disrupt data, we expect the quality of data to vary by task and situation. Put differently, the microstructure of the strategic environment is very important. We expect data about friendly forces to be more reliable because commanders can mandate reporting formats and schedules. We expect logistics and administration reports to be more reliable than combat reporting, which is more exposed to enemy interaction. Intelligence about enemy dispositions and capabilities should be even less reliable. Even so, intelligence about “puzzles” (such as the locations and capabilities of weapon systems) may be more reliable than intelligence about “mysteries” (such as future intentions and national resolve).⁴⁴ Data from technical sensors tend to be better structured and more voluminous than human intelligence reports, which require significant interpretation. Enemy deception or disinformation operations tend to undermine data quality, as does intelligence politicization for parochial interests. It is critical to assess the specific strategic context of data, and thus the suitability of AI, for any given decision task.

41. *Ibid.*, p. 119.

42. *Inter alia*, James A. Russell, *Innovation, Transformation, and War: Counterinsurgency Operations in Anbar and Ninewa Provinces, Iraq, 2005–2007* (Stanford, Calif.: Stanford University Press, 2010); Timothy S. Wolters, *Information at Sea: Shipboard Command and Control in the U.S. Navy, from Mobile Bay to Okinawa* (Baltimore, Md.: Johns Hopkins University Press, 2013); and Nina A. Kollars, “War’s Horizon: Soldier-Led Adaptation in Iraq and Vietnam,” *Journal of Strategic Studies*, Vol. 38, No. 4 (2015), pp. 529–553, <https://doi.org/10.1080/01402390.2014.971947>.

43. Generally, Lindsay, *Information Technology and Military Power*.

44. Gregory F. Treverton, *Reshaping National Intelligence for an Age of Information* (New York: Cambridge University Press, 2003), pp. 11–13.

JUDGMENT AND MILITARY INSTITUTIONS

Even if there are enough of the right type of data, AI still relies on people to determine what to predict and why. Commercial firms, for instance, make many different judgments in determining their business models, corporate values, labor relations, and negotiating objectives. Military organizations face analogous management challenges, but they also face unique ones. Military judgment also encompasses national interests, political preferences, strategic missions, commander's intent, rules of engagement, combat ethics, and martial socialization. Because the costs and consequences of war are so profound, all these problems tend to be marked by ambiguity, controversy, and painful trade-offs. Judgment thus becomes more difficult, and ever more consequential, in military affairs.

There are three types of machine learning algorithms.⁴⁵ All require human judgment. First, in "supervised learning," the human tells the machine what to predict. Second, "unsupervised learning" requires judgment about what to classify and what to do with the classifications. Third, "reinforcement learning" requires advance specification of a reward function. A reward function assigns a numerical score to the perceived state of the world to enable a machine to maximize a goal. More complicated strategies may combine these approaches by establishing instrumental goals in pursuit of the main objective. In every case, a human ultimately codes the algorithm and defines the payoffs for the machine.

In economic terms, judgment is the specification of the utility function.⁴⁶ The preferences and valuations that determine utility are distinct from the strategies that maximize it. To take a trivial example, people who do not mind getting wet and dislike carrying umbrellas will not carry one, regardless of the weather forecast. People who dislike getting wet and do not mind carrying umbrellas might always have an umbrella in their bag. Others might carry an umbrella if the chance of rain is 75 percent but not if it is 25 percent. The prediction of rain is independent of preferences about getting wet or being prepared to get wet. Similarly, the AI variation on the notorious "trolley problem" poses an ethical dilemma about life-or-death choices. For example, should a self-driving car swerve to avoid running over four children at the risk

45. For an accessible introduction, see Sejnowski, *The Deep Learning Revolution*.

46. Critiques of economic rationality that appeal to cognitive psychology or social institutions underscore the importance of judgment. See Rose McDermott, *Political Psychology in International Relations* (Ann Arbor: University of Michigan Press, 2004); and Janice Gross Stein, "The Micro-Foundations of International Relations Theory: Psychology and Behavioral Economics," *International Organization*, Vol. 71, No. S1 (2017), pp. S249–S263, <https://doi.org/10.1017/S0020818316000436>.

of killing its human passenger? If the AI predicts even odds that someone will die either way, the car should swerve if all lives are equally valuable, but it should not swerve if the passenger's life is worth at least four times as much as that of a random child. This somewhat contrived dilemma understates the complexity of the judgment involved. Indeed, the ethical dilemmas of AI reinvigorate longstanding critiques of utilitarian reasoning. As Heather Roff points out, "We cannot speak about ethical AI because all AI is based on empirical observations; we cannot get an 'ought' from an 'is.' If we are clear eyed about how we build, design, and deploy AI, we will conclude that all of the normative questions surrounding its development and deployment are those that humans have posed for millennia."⁴⁷

If the trolley problem seems far-fetched, consider the case of a self-driving Uber car that killed a cyclist in Tempe, Arizona.⁴⁸ The AI had predicted a low but nonzero probability that a human was in its path. The car was designed with a threshold for ignoring low-probability risks. The priority of not hitting humans was obvious enough. Yet, with an error tolerance set to zero, the car would not be able to drive. The question of where to set the tolerance was a judgment call. In this case, it appears that the prespecified judgment was tragically inappropriate for the context, but the prediction machine had absolutely no concept of what was at stake.

A well-specified AI utility function has two characteristics. First, goals are clearly defined in advance. If designers cannot formally specify payoffs and priorities for all situations, then each prediction will require a customized judgment. This is often the case in medical applications.⁴⁹ When there are many possible situations, human judgment is often needed upon seeing the diagnosis. The judgment cannot be determined in advance because it would take too much time to specify all possible contingencies. Such dynamic or nuanced situations require, in effect, incomplete contracts that leave out complex, situation-specific details to be negotiated later.⁵⁰ Because all situations cannot be stipulated in advance, judgment is needed after seeing the prediction to interpret the spirit of the agreement.

47. Heather M. Roff, *The Folly of Trolleys: Ethical Challenges and Autonomous Vehicles* (Washington, D.C.: Brookings Institution Press, December 17, 2018), <https://www.brookings.edu/research/the-folly-of-trolleys-ethical-challenges-and-autonomous-vehicles/>.

48. National Transportation Safety Board [NTSB], "Highway Accident Report: Collision between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018," Highway Accident Report NTSB/HAR-19/03 (Washington, D.C.: NTSB, November 19, 2019), <https://trid.trb.org/view/1751168>.

49. Trevor Jamieson and Avi Goldfarb, "Clinical Considerations When Applying Machine Learning to Decision-Support Tasks versus Automation," *BMJ Quality & Safety*, Vol. 28, No. 10 (2019), pp. 778–781, <https://doi.org/10.1136/bmjqs-2019-009514>.

50. Williamson, "The Economics of Organization."

The military version of incomplete contracting is “mission command,” which specifies the military objective and rules of engagement but empowers local personnel to interpret guidance, coordinate support, and tailor operations as the situation develops.⁵¹ The opposite of mission command, sometimes described as “task orders,” is more like a complete contract that tells a unit exactly what to do and how to do it. Standard operating procedures, doctrinal templates, and explicit protocols help to improve the predictability of operations by detailing instructions for operations and equipment handling. In turbulent environments with unpredictable adversaries, however, standardized task orders may be inappropriate. The greater the potential for uncertainty and accident in military operations, the greater the need for local commanders to exercise initiative and discretion. In Clausewitzian terms, “fog” on the battlefield and “friction” in the organization require commanders to exercise “genius,” which is “a power of judgment raised to a marvelous pitch of vision, which easily grasps and dismisses a thousand remote possibilities which an ordinary mind would labor to identify and wear itself out in so doing.”⁵² The role of “genius” in mission command becomes particularly important, and particularly challenging, in modern combined arms warfare and multi-domain operations.⁵³ When all possible combinations of factors cannot possibly be specified in advance, personnel have to exercise creativity and initiative in the field. Modern military operations tend to mix elements of both styles by giving local commanders latitude in how they interpret, implement, and combine the tools, tactics, and procedures that have been standardized, institutionalized, and exercised in advance.

The second characteristic of a well-specified AI utility function is that all stakeholders should agree on what goals to pursue. When it is difficult for people to agree on what to optimize, transparent institutional processes for evaluating or aggregating different preferences may help to validate or legitimate decisions that guide AI systems. Unfortunately, consensus becomes elusive as “genius” becomes more geographically distributed, socially collaborative, and technically exacting.⁵⁴ In an ethnography of divisional command in

51. For example, Department of the Army, *ADP 6-0: Mission Command: Command and Control of Army Forces*, Army Doctrine Publication No. 6-0 (Washington, D.C.: U.S. Department of the Army, May 17, 2012), https://armypubs.army.mil/epubs/DR_pubs/DR_a/ARN18314-ADP_6-0-000-WEB-3.pdf.

52. Clausewitz, *On War*, p. 112.

53. Biddle, *Military Power*; Bart Van Bezooijen and Eric-Hans Kramer, “Mission Command in the Information Age: A Normal Accidents Perspective on Networked Military Operations,” *Journal of Strategic Studies*, Vol. 38, No. 4 (2015), pp. 445–466, <https://doi.org/10.1080/01402390.2013.844127>; and Ryan Grauer, *Commanding Military Power: Organizing for Victory and Defeat on the Battlefield* (Cambridge: Cambridge University Press, 2016).

54. John Ferris and Michael I. Handel, “Clausewitz, Intelligence, Uncertainty, and the Art of Com-

Afghanistan, Anthony King writes that “a general must define a mission, manage the tasks of which it is comprised and motivate the troops.”⁵⁵ The first of these three factors—specifying positive objectives and negative limitations—is the consummate function of judgment; AI offers little help here. AI might provide some support for the second factor, oversight and administration, which involves a mixture of judgment and prediction. The third factor is leadership, which is fundamentally a judgment problem insofar as leaders attempt to socialize common purposes, values, and interpretations throughout an organization. Again, AI is of little use for issues of leadership, which become more important as organizations become geographically and functionally distributed: “The bureaucratic expertise of the staff has been improved and their cohesiveness has been condensed so that they are now bound in dense solidarity, even when they are not co-present.”⁵⁶ Indeed, “decision-making has proliferated” in all three areas—strategy, management, leadership—because a “commander can no longer direct operations alone.”⁵⁷ According to King, the commanding general is now less of a central controller and more of a social focal point for coordinating the complex interactions of “the command collective.”

Collective command, however, is a collective action problem. In some cases, standard operating procedures and socialization rituals can simplify judgment tasks. King finds that “command teams, command boards, principal planning groups and deputies have appeared to assist and to support the commander and to manage discrete decision cycles to which the commander cannot attend.”⁵⁸ Yet, in other cases, personnel from different services, branches, or units may disagree over how to interpret even basic tactics, techniques, and procedures.⁵⁹ Disagreement may turn into controversy when mission assignments fall outside the scope of what professionals deem right or appropriate, as when armies are tasked with counterinsurgency, air forces are tasked with close air support, or cultural preferences clash.⁶⁰ More serious disagreements

mand in Military Operations,” *Intelligence and National Security*, Vol. 10, No. 1 (1995), pp. 1–58, <https://doi.org/10.1080/02684529508432286>.

55. Anthony King, *Command: The Twenty-First-Century General* (Cambridge: Cambridge University Press, 2019), p. 438.

56. *Ibid.*, p. 443.

57. *Ibid.*, p. 439.

58. *Ibid.*, p. 440.

59. Harvey M. Sapolsky, Eugene Gholz, and Caitlin Talmadge, *US Defense Politics: The Origins of Security Policy*, 3rd ed. (New York: Routledge, 2017), pp. 93–116.

60. On the cultural origins of military preferences, see Elizabeth Kier, *Imagining War: French and British Military Doctrine between the Wars* (Princeton, N.J.: Princeton University Press, 1997); Jeffrey W. Legro, *Cooperation under Fire: Anglo-German Restraint during World War II* (Ithaca, N.Y.: Cornell University Press, 2013 [1995]); and Austin Long, *The Soul of Armies: Counterinsurgency Doctrine and Military Culture in the US and UK* (Ithaca, N.Y.: Cornell University Press, 2016).

about war aims and military methods can emerge within the civil-military chain of command or among coalition partners.⁶¹ Just as data availability and bias vary for any given decision task, we also expect variability in the clarity and consensus of judgment. Any factors that exacerbate confusion or disagreement in military institutions should be expected to make judgment more difficult for AI automation.

AI Performance in Military Decision-Making Tasks

As we have explained in the previous sections, decision-making is a universal process, but decision inputs are context specific. Even if the same AI technology is available to all organizations, the strategic and institutional conditions that have enabled AI success in the business world may not be present in war. We thus infer two general hypotheses about the key AI complements of data and judgment. First, stable, cooperative environments are more conducive for plentiful, unbiased data; conversely, turbulent, competitive environments tend to produce limited, biased data. Second, institutional standardization and solidarity encourage well-defined, consensual judgments; conversely, idiosyncratic local practices and internal conflict lead to ambiguous, controversial judgments.

The combination of these hypotheses describes four different regimes of AI performance in military decision-making tasks. Table 2 summarizes these categories by synthesizing the strategic and institutional inputs from the decision-making context (figure 1) into the human-machine division of labor (table 1) for key military functions. The best case for AI performance is what we call “automated decision-making.” Quality data and clear judgment are most likely to be available in highly routinized administrative and logistics tasks that are more analogous to civilian organizational tasks. Anything that bureaucracies can do well, AI can probably help them to do better. The worst case for AI is the opposite quadrant, in which both automation complements are absent. We label this category “human decision-making” because AI cannot perform tasks characterized by limited, biased data and ambiguous, controversial judgments. For military strategy and command tasks, the Clausewitzian extremes of “fog” in the environment and “friction” in the organization require human “genius.” In the other two quadrants, in which one necessary complement is present but the other is absent, the human-machine division of labor is

61. For example, Risa Brooks, *Shaping Strategy: The Civil-Military Politics of Strategic Assessment* (Princeton, N.J.: Princeton University Press, 2008); and Jeremy Pressman, *Warring Friends: Alliance Restraint in International Politics* (Ithaca, N.Y.: Cornell University Press, 2012).

Table 2. How Strategic and Institutional Conditions Shape AI Performance in Military Tasks

		Effect of Environment on Data	
		<i>stability and cooperation</i> → <i>high-quality data</i>	<i>turbulence and competition</i> → <i>low-quality data</i>
Effect of Institutions on Judgment	<i>standardization and solidarity</i> → <i>clear judgment</i>	Automated Decision-Making Full automation can increase the scale and efficiency of highly bureaucratized administrative and logistics tasks.	Premature Automation Full automation in complex fire and maneuver tasks heightens the risks of targeting error and inadvertent escalation.
	<i>idiosyncrasy and conflict</i> → <i>difficult judgment</i>	Human-Machine Teaming Automated decision aids for intelligence analysis and operational planning can augment, but not replace, human decision-making.	Human Decision-Making Strategy and command will continue to rely on human interpretation and leadership.

more complicated. The category of “premature automation” describes situations in which machines receive clear goals in uncertain environments. If rapid, automated decisions are tightly coupled to lethal actions, this is a recipe for disaster. The mismatch between the evolving situation and the decisions encoded in AI systems may not be immediately obvious to humans, which heightens the risks of tragic outcomes such as targeting error or inadvertent escalation. In the converse category of “human-machine teaming,” judgment tasks are difficult, but quality data are available. AI decision aids (e.g., graphs, tables, map overlays, image annotations, and textual summaries) provide predictions that augment but do not replace human decisions, while humans maintain a healthy skepticism about AI shortcomings. Many operational planning and intelligence analysis tasks fall into this category, along with tactical mission support tools. To demonstrate the plausibility of our framework, we next offer a few commercial examples and explore potential military applications and further implications in each category.

AUTOMATED DECISION-MAKING

Full automation of decision-making can improve performance if there are quality data and clear judgments. For example, Australia’s Pilbara region has large quantities of iron ore. The mining sites are far from any major city, and the local conditions are often so hot that it is hazardous for humans to

work there. Since 2016, mining giant Rio Tinto has deployed dozens of self-driving trucks.⁶² These trucks have saved operating costs while reducing risk to human operators. Such automation is feasible because the data are plentiful relative to the needs at hand—the trucks drive on the same roads each day, and there are few surprises in terms of human activity. Data collection is therefore limited to a small number of roads with few obstacles. The main task for the AI is to predict whether the path is clear. Once this prediction is made, the judgment is well defined and easy to specify in advance: if the path is clear, continue; if it is not clear, stop and wait. In other examples of successful automation, robotic cameras have been deployed in a variety of settings, including during basketball games, swimming competitions, and “follow me” aerial drones.⁶³

We expect AI to be useful for military tasks with clear civilian analogs. While much of the popular debate about military AI is preoccupied with automated weaponry, it is likely that many promising applications will be designed to support bureaucratic functions. Bureaucracies are, among other things, computational systems that gather and process data to render operations more legible, predictable, and controllable.⁶⁴ The most bureaucratized parts of a military organization are therefore good candidates for computational automation. Administrative transactions tend to be repetitious, which generates a large amount of high-quality data that can be used for training and prediction. Organizations are attracted to standardization because it makes their equipment, procedures, and personnel easier to count, compare, and control.⁶⁵ Procedural standardization also constrains organizational behavior, which makes it easier for managers to specify judgments in advance. Moreover, peacetime administrative tasks are somewhat less exposed to battlefield turbulence, reducing the requirement for last-minute interpretation.

We expect automation to improve the efficiency and scale of routinized activities that entail filling in missing information, measuring technical performance, tracking personnel, and anticipating future needs. Indeed, AI may enhance many routine tasks associated with developing budgets, recruiting and training personnel, identifying leadership potential, scheduling unit rosters, designing and procuring weapon systems, planning and evaluating exer-

62. Agrawal, Gans, and Goldfarb, *Prediction Machines*, pp. 113–114.

63. *Ibid.*, p. 115.

64. James G. March and Herbert A. Simon, *Organizations* (New York: John Wiley, 1958).

65. JoAnne Yates, *Control through Communication: The Rise of System in American Management* (Baltimore, Md.: Johns Hopkins University Press, 1989); and Wendy Nelson Espeland and Mitchell L. Stevens, “Commensuration as a Social Process,” *Annual Review of Sociology*, Vol. 24, No. 1 (1998), pp. 313–343, <https://doi.org/10.1146/annurev.soc.24.1.313>.

cises, caring for the morale and welfare of service members and their families, and providing health care to service members and veterans.⁶⁶ At the same time, it is important to recognize that seemingly trivial procedures can become politicized when budgets and authorities are implicated.⁶⁷ Even in the absence of parochialism, the complexity of administrative systems introduces interpretive challenges for personnel. These internal frictions undermine the conditions for successful administrative automation.

Logistics supply chains may also be good candidates for automation. Indeed, firms like DHL and FedEx have leveraged AI to streamline their delivery networks. Standardized parts, consumption rates, repetitive transactions, and preventive maintenance schedules generate abundant data about defined tasks. Using historical performance data, predictive maintenance systems can monitor consumption rates and automatically order replacement parts before a weapon or platform breaks. For example, one U.S. Air Force system uses a predictive algorithm to decide when mechanics should perform an inspection, which allows them to tailor the maintenance and repairs for individual aircraft rather than adhere to generic schedules.⁶⁸ But we contend that the prediction of supply and demand for just-in-time delivery will be more difficult in war. While bureaucrats may be insulated from the turmoil of the battlefield, supply lines are more exposed. The enemy can interdict or sabotage logistics. As wartime attrition consumes spare parts, units may squabble about which ones should be resupplied. Friendly units may resort to using platforms and parts in unconventional ways. All this turbulence will cause predictions to fail, which essentially shifts AI into the category of premature automation, discussed below. The classical military solution to such problems is to stockpile an excess of supplies, precisely because wartime consumption is so hard to predict.⁶⁹ If organizations eliminate slack resources with AI systems in pursuit of efficiency, however, then they may sacrifice effectiveness when systems encounter unforeseen circumstances.⁷⁰

66. Brian David Ray, Jeanne F. Forgey, and Benjamin N. Mathias, "Harnessing Artificial Intelligence and Autonomous Systems across the Seven Joint Functions," *Joint Force Quarterly*, Vol. 96, No. 1 (2020), p. 124, <https://ndupress.ndu.edu/Portals/68/Documents/jfq/jfq-96/jfq-96.pdf>; and Stephan De Spiegeleire, Matthijs Maas, and Tim Sweijts, *Artificial Intelligence and the Future of Defense: Strategic Implications for Small- and Medium-Sized Force Providers* (The Hague, Netherlands: Hague Centre for Strategic Studies, 2017), pp. 91–94, <https://www.jstor.org/stable/resrep12564.1>.

67. Sapolsky, Gholz, and Talmadge, *US Defense Politics*.

68. Stoney Trent and Scott Lathrop, "A Primer on Artificial Intelligence for Military Leaders," *Small Wars Journal*, August 22, 2018, <http://smallwarsjournal.com/jrnl/art/primer-artificial-intelligence-military-leaders>.

69. Martin van Creveld, *Supplying War: Logistics from Wallenstein to Patton*, 2nd ed. (New York: Cambridge University Press, 2004).

70. Loss of resilience is a longstanding concern associated with organizational automation. See

In sum, we expect AI to be most useful for automating routine tasks that are bureaucratically insulated from battlefield turbulence. Administration and logistics tasks that are repetitious and standardized are more likely to have both quality data and clear goals. Humans still provide judgment to define those clear goals, but this happens in advance. Although these conditions are ideal for automation, they can be elusive in practice, especially if there are contested resources and personnel decisions. As a result, even the low-hanging fruit applications will often fall into the other three categories in table 2, particularly human-machine teaming.

HUMAN DECISION-MAKING

At the other extreme, humans still make all the decisions for situations in which data are of low quality and judgment is difficult. Machine predictions degrade without quality data. Fortunately, because judgment is also difficult in this category, there is little temptation to automate. There are no commercial examples in this category because we have not seen AI systems that successfully found companies, lead political movements, or set legal precedents by themselves. Without quality data and clear judgment, such machines would be of little use. As long as advances in machine learning are best understood as improvements in prediction rather than AGI, then tasks in this category will require human beings. People do not always make good decisions, however, and successful decisions depend on many different psychological and social factors, plus good luck.

Strategy abounds with complex and controversial political and moral judgments. What is worth fighting for, or compromising on? When should allies be embraced, or abandoned? When is butter worth more than guns, and when does the stability of deterrence outweigh the pursuit of power? For what national interests should men and women be inspired to kill, and die? And when should killers show restraint? The answers to these questions spring from many sources such as ideology, psychology, and domestic politics, but they do not come from machines. AI systems may win games like *Jeopardy*, *Go*, and complex video games. War shares some features with some games, such as strategic competition and zero-sum payoffs. But war is not a game. Games are defined by institutionalized rules, but the failure of institutions in anarchy gives rise to war.

In Clausewitzian terms, war is the use of violence to impose one's will on a reactive opponent. The interaction of political rationality, national passion, and

random chance gives war a chaotic quality.⁷¹ The problems of formulating political-military strategy and commanding forces in battle live in the heart of this chaos. Curiosity, creativity, grit, and perseverance become important character traits, to say nothing of empathy, mercy, and compassion. Whenever “fog” and “friction” are greatest, and human “genius” is most in demand, there is little role for AI. Humans will often fail in these circumstances, too, but at least they have a fighting chance.

AI systems may still be able to support strategy and command by providing decision aids that improve the intelligence, planning, and administrative inputs to decision-making. Yet, this simply underscores the importance of decomposing decision tasks into subtasks that AI can support and subtasks that humans must perform. The partition of judgment tasks itself is an act of judgment. For fluid, ambiguous, or controversial practices, which are common in strategy and command, the boundaries of data, judgment, prediction, and action may be difficult to distinguish from each other, let alone from other decision tasks. Judgment becomes even more important in these situations.

PREMATURE AUTOMATION

In between the extremes of fully automated and fully human decision-making, there are both risks and opportunities. The mixed cases of premature automation and human-machine teaming generate most of the worry and excitement about AI. Reliance on AI is particularly risky in situations in which the data are of low quality, but the machine is given clear objectives and authorized to act. The risks are greatest when lethal action is authorized. If data are biased or incomplete, then it would be better to let humans rather than machines interpret the evolving situation (i.e., human decision-making). If humans mistakenly believe that data are abundant and unbiased, however, then they may wrongly assume that critical tasks can be delegated to AI (i.e., automated decision-making). Automation seems seductively feasible, but the definition of the reward function fails to keep up with important changes in the context of decision-making. Clear judgment in a turbulent environment creates hidden dangers.

An example of premature automation is when Amazon built an AI system to help with its hiring processes.⁷² The judgment seemed clear: the machine should select the workers who are likely to succeed in the company. Amazon

71. Alan Beyerchen, “Clausewitz, Nonlinearity, and the Unpredictability of War,” *International Security*, Vol. 17, No. 3 (Winter 1992/93), pp. 59–90, <https://doi.org/10.2307/2539130>.

72. Jeffrey Dastin, “Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women,” Reuters, October 10, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>.

receives thousands of résumés, and a better screening tool could automate the many hours that human recruiters spend screening them. There were reasons to be optimistic that this automation would reduce bias and yield higher quality and more diverse candidates.⁷³ Unfortunately, Amazon's past applications and hiring practices meant that the data contained insufficient examples of successful women applicants. Without data on successful women applicants, the AI learned that Amazon should not hire women, and it consequently screened out résumés that included the word "women." The existing biases in organizational processes produced biased AI training data. Fortunately, Amazon management realized that the AI exacerbated rather than solved Amazon's existing problems, and the company never deployed this AI tool.

Data may also be biased because they are based on decisions that a prediction machine may not understand. For example, an early AI for chess was trained on thousands of Grandmaster games. When deployed, the program sacrificed its queen early in the game because it had learned that Grandmasters who do so tend to win. Human Grandmasters only sacrifice their queen, however, when doing so generates a clear path to victory.⁷⁴ While this issue has been solved in AI chess, the underlying challenge continues. Even when the utility function is clear, training data is often a result of tacit assumptions in human decision-making. Sometimes those human decisions—such as sacrificing a queen in chess—create biased data that cause AI predictions to fail.

In contrast with chess, the risks of premature automation are more extreme in the military realm (e.g., fratricide and civilian casualties), but the logic is the same. Militaries abound with standard operating procedures and tactical doctrines that guide the use of lethal capabilities (e.g., instructions for the safe operation of weapon platforms, playbooks for tactical maneuvers, and policies for employing weapons). To the degree that goals and mission parameters can be clearly specified, tactical operations will appear to be attractive candidates for automation. To the degree that combat timelines are expected to be extremely compressed, moreover, automation may appear to be even more urgent.⁷⁵ Rapid decision-making would necessitate the pre-specification of goals and payoffs and the coupling of AI prediction to robotic action. Lethal autonomous weapon systems use prediction to navigate complex environments in order to arrive at destinations or follow targets, within constraints that are supplied by human operators.⁷⁶ Their targeting systems base their predictions

73. Kleinberg et al., "Discrimination in the Age of Algorithms."

74. Agrawal, Gans, and Goldfarb, *Prediction Machines*, p. 63.

75. Horowitz, "When Speed Kills."

76. Heather M. Roff and Richard Moyes, "Meaningful Human Control, Artificial Intelligence, and

on training data that identify valid targets. Using algorithms, machines may rapidly and accurately identify targets at far greater distances than human visual recognition, and algorithmic target recognition may be collocated with sensors to reduce response times.⁷⁷

Many AI-enabled weapons already or imminently exist. The Israeli Harpy loitering munition can search for and automatically engage targets, and China has plans for similar “intelligentized” cruise missiles.⁷⁸ Russia is developing a variety of armed, unmanned vehicles capable of autonomous fire or unarmed mine clearance.⁷⁹ The United States has been exploring combat applications for AI in all warfighting domains. In the air, the “Loyal Wingman” program pairs an unmanned F-16 with a manned F-35 or F-22 to explore the feasibility of using humans to direct autonomous aircraft, such as the XQ-58A Valkyrie.⁸⁰ Air combat algorithms that can process sensor data and plan effective combat maneuvers in the span of milliseconds have already defeated human pilots in some simulators.⁸¹ At sea, the U.S. Navy’s LOCUST project explores the feasibility of launching swarms of expendable surface-to-air drones.⁸² The Defense Advanced Research Projects Agency’s (DARPA) Continuous Trail Unmanned Vessel program is designed to search for enemy missile submarines and automatically trail them for months at a time, reporting regularly on their locations.⁸³ On land, U.S. Marine “warbot companies” equipped with networks of small robots might provide distributed sensing and precision fire.⁸⁴ Auto-

Autonomous Weapons,” paper prepared for the Informal Meeting of Experts on Lethal Autonomous Weapons Systems, UN Convention on Certain Conventional Weapons, Geneva, April 11–15, 2016, <https://article36.org/wp-content/uploads/2016/04/MHC-AI-and-AWS-FINAL.pdf>.

77. See, for example, the Defense Advanced Research Projects Agency’s (DARPA) Target Recognition and Adaption in Contested Environments (TRACE) program, described in John Keller, “DARPA TRACE program using advanced algorithms, embedded computing for radar target recognition,” *Military and Aerospace Electronics*, July 23, 2015, <https://www.militaryaerospace.com/computers/article/16714226/darpa-trace-program-using-advanced-algorithms-embedded-computing-for-radar-target-recognition>.

78. Elsa B. Kania, *Battlefield Singularity: Artificial Intelligence, Military Revolution, and China’s Future Military Power* (Washington, D.C.: Center for a New American Security, November 2017), <https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/Battlefield-Singularity-November-2017.pdf>.

79. Spiegeleire, Maas, and Sweijts, *Artificial Intelligence and the Future of Defense*, pp. 80, 82.

80. Hoadley and Saylor, *Artificial Intelligence and National Security*, p. 13.

81. Norine MacDonald and George Howell, “Killing Me Softly: Competition in Artificial Intelligence and Unmanned Aerial Vehicles,” *PRISM*, Vol. 8, No. 3 (2019), pp. 103–126, https://ndupress.ndu.edu/Portals/68/Documents/prism/prism_8-3/prism_8-3.pdf.

82. LOCUST stands for “Low-Cost Unmanned Aerial Vehicle Swarming Technology.” See Jules Hurst, “Robotic Swarms in Offensive Maneuver,” *Joint Force Quarterly*, Vol. 87, No. 4 (2017), pp. 105–111, https://ndupress.ndu.edu/Portals/68/Documents/jfq/jfq-87/jfq-87_105-111_Hurst.pdf?ver=2017-09-28-093018-793.

83. Spiegeleire, Maas, and Sweijts, *Artificial Intelligence and the Future of Defense*, p. 91.

84. Jeff Cummings et al., “Marine Warbot Companies: Where Naval Warfare, the U.S. National Defense Strategy, and Close Combat Lethality Task Force Intersect,” *War on the Rocks*, June 28,

mated counter-battery responses, which accurately retaliate against the origin of an attack, could give human commanders leeway to focus on second- and third-order decisions in the wake of an attack.⁸⁵ In the cyber domain, AI systems might autonomously learn from and counter cyberattacks as they evolve in real time, as suggested by the performance of the Mayhem system in DARPA's 2016 Cyber Grand Challenge.⁸⁶ AI could be especially useful for detecting new signals in the electromagnetic spectrum and reconfiguring electronic warfare systems to exploit or counter them.⁸⁷ Space satellites, meanwhile, have been automated from their inception, and space operations might further leverage AI to enhance surveillance and control.

Much of the AI security literature is preoccupied with the risks posed by automated weapons to strategic stability and human security.⁸⁸ Risks of miscalculation will increase as the operational context deviates from the training data set in important or subtle ways. The risk of deviation increases with the complexity and competitiveness of the strategic environment, while the costs of miscalculation increase with the lethality of automated action. The machine tries to optimize a specific goal, but in the wrong context, doing so can lead to false positives. AI weapons may inadvertently either target innocent civilians or friendly forces or trigger hostile retaliation. In these cases, the AI would have the authority to kill but would not understand the ramifications. The risks are particularly stark in the nuclear arena. Nuclear war is the rarest of rare events—keeping it that way is the whole point of nuclear deterrence—so training data for AI systems is either nonexistent or synthetic (i.e., based on simulation).⁸⁹ Any tendency for AI systems to misperceive or miscalculate when confronted with uncertain or novel situations could have catastrophic consequences.⁹⁰ In short, autonomous weapon systems that combine prediction with action can quickly make tragic mistakes.

2018, <https://warontherocks.com/2018/06/marine-warbot-companies-where-naval-warfare-the-u-s-national-defense-strategy-and-close-combat-lethality-task-force-intersect/>.

85. Ray, Forgey, and Mathias, "Harnessing Artificial Intelligence and Autonomous Systems across the Seven Joint Functions," p. 123.

86. Spiegeleire, Maas, and Sweijs, *Artificial Intelligence and the Future of Defense*, p. 88.

87. Matthew J. Florenzen, Kurt M. Shulkitas, and Kyle P. Bair, "Unmasking the Spectrum with Artificial Intelligence," *Joint Force Quarterly*, Vol. 95, No. 4 (2019), pp. 116–123, https://ndupress.ndu.edu/Portals/68/Documents/jfq/jfq-95/jfq-95_116-123_Florenzen-Skulkitas-Bair.pdf.

88. For example, Payne, "Artificial Intelligence"; and Suchman, "Algorithmic Warfare and the Reinvention of Accuracy."

89. Rafael Loss and Joseph Johnson, "Will Artificial Intelligence Imperil Nuclear Deterrence?" *War on the Rocks*, September 19, 2019, <https://warontherocks.com/2019/09/will-artificial-intelligence-imperil-nuclear-deterrence/>.

90. Vincent Boulanin, ed., *The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk*, Vol. 1, *Euro-Atlantic Perspectives* (Solna, Sweden: Stockholm International Peace Research Institute, May 2019), <https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability->

For most of the examples reviewed in this section, humans should adjust goals on a case-by-case basis to avoid the substantial operational risks of full automation. The DARPA Air Combat Evolution (ACE) program, which trains AI pilots in dogfight simulations, highlights risks that can emerge when AI is given too much decision autonomy in rapidly changing contexts: “at one point in the AlphaDogfight trials, the organisers threw in a cruise missile to see what would happen. Cruise missiles follow preordained flight paths, so they behave more simply than piloted jets. The AI pilots struggled with this because, paradoxically, they had beaten the missile in an earlier round and were now trained for more demanding threats.”⁹¹ Experiments like this have encouraged ACE to focus on “manned-unmanned teaming” rather than full autonomy. The engineering challenge is then to partition the cognitive load of prediction and judgment correctly (i.e., to decompose the task into different subtasks) so that faster machines and mindful humans can play to their strengths. These examples show that automation risk from low-quality data increases the importance of careful human judgment and teaming a human with the machine. Human personnel are needed to identify when data are incomplete or biased in the specific context of any given decision, and to provide judgment on how to act on potentially inaccurate predictions. For many tactical fire and maneuver tasks, full automation is prohibitively risky, but close human supervision may be able to mitigate that risk.

HUMAN-MACHINE TEAMING

If quality data are available but judgment is difficult, then AI can still provide predictions if humans first tell the machines what to do. We describe this category as “human-machine teaming” because skilled people can use AI to enhance decision-making, but they must guide and audit AI performance in sensitive or idiosyncratic circumstances. In these situations, quality data will generate reliable predictions. Owing to the difficulty of prespecifying judgment, however, most practitioners are not tempted to deploy full automation because they recognize that doing so may risk creating more bad decisions.

Consider the civilian example of tax law and the ambiguity about whether investment income should be taxed as business income or capital gains. Typically, a company would hire a lawyer to collect facts on the case and predict what the courts are likely to find. Then, the lawyer would advise the client

nuclear-risk.pdf; and Lora Saalman, “Fear of False Negatives: AI and China’s Nuclear Posture,” *Bulletin of the Atomic Scientists* blog, April 24, 2018, <https://thebulletin.org/2018/04/fear-of-false-negatives-ai-and-chinas-nuclear-posture/>.

91. “Fighter Aircraft Will Soon Get AI Pilots,” *Economist*, November 19, 2020, <https://www.economist.com/science-and-technology/2020/11/15/fighter-aircraft-will-soon-get-ai-pilots>.

on a course of action. One firm developed an AI that scans tax law decisions to predict tax liabilities. The AI does not recommend a course of action because making that judgment requires knowing the client's risk preferences and comfort navigating the legal system. The AI predicts what would happen if the case were to go to court, but it cannot determine whether going to court is a good idea. Legal decisions in this task are the product of human-machine teaming between the predictive AI and the human lawyer, who must interpret the prediction to judge what advice best serves the client.⁹²

For similar reasons, human-machine teaming ought to flourish in intelligence and planning organizations. There is much growth potential for AI in units that are awash in intelligence, surveillance, and reconnaissance (ISR) data.⁹³ In remotely piloted aircraft (drone) operations, for instance, the information processing burden is intense at even the most tactical level. According to an ethnographic study by Timothy Cullen, "to fly the aircraft and control the sensor ball, Reaper and Predator crews had to coordinate the meaning, movement, and presentation of a myriad of menus, windows, and tables on 16 displays and 4 touch screens with 4 keyboards, 2 trackballs, 2 joysticks, and 8 levers."⁹⁴ Cullen describes a complicated mixture of prediction and judgment as "operators negotiated and constructed a constrained environment in the ground control station to coordinate verbal, typed, written, pictorial, and geographical representations of a mission; to identify patterns in scenes from the aircraft's sensors; and to associate those patterns with friendly and enemy activity."⁹⁵ ISR drones generate so much data, "37 years of full motion footage in 2011 alone," that "much of the collect goes unanalyzed."⁹⁶ AI is able to alleviate some of the data processing burden by continuously monitoring multiple data feeds and highlighting patterns of interest. Yet, Cullen highlights ways in which aircrew make many, seemingly minor, value judgments about what they should—and should not—be doing with their sensors and weapons. In other words, AI provides one complement to the data—the prediction—but it does not provide the judgment that also underlies decision-making.

In principle, many intelligence tasks might benefit from machine learning.

92. Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb, "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction," *Journal of Economic Perspectives*, Vol. 33, No. 2 (Spring 2019), p. 35, <https://doi.org/10.1257/jep.33.2.31>.

93. Keith Dear, "A Very British AI Revolution in Intelligence Is Needed," *War on the Rocks*, October 19, 2018, <https://warontherocks.com/2018/10/a-very-british-ai-revolution-in-intelligence-is-needed/>.

94. Timothy M. Cullen, "The MQ-9 Reaper Remotely Piloted Aircraft: Humans and Machines in Action," Ph.D. dissertation, Massachusetts Institute of Technology, 2011, p. 272.

95. *Ibid.*, p. 273.

96. Dear, "A Very British AI Revolution in Intelligence Is Needed."

Image recognition algorithms can sift through drone video feeds to identify enemy activity. Facial recognition systems can detect individual targets of interest, while emotional prediction algorithms can aid in identifying the hostile or benign intent of individuals on a crowded street. Speech recognition, voice synthesis, and translation systems can alleviate shortages of human translators for human and signals intelligence, as well as for civil-military relations and information operations. In general, AI is well suited to the intelligence task of analyzing bulk data and identifying patterns, for example, in identifying and tracking terrorist groups or insurgents.⁹⁷

In practice, intelligence is often more art than science. Intelligence professionals deal with deceptive targets, ambiguous data, and subtle interpretations.⁹⁸ Unforeseen changes in the strategic environment or mission objectives create new data requirements or, worse, undermine the referential integrity of existing data. On a case-by-case basis, practitioners draw on their subject matter expertise and experience to make judgments. Applying this judgment to an AI prediction is a difficult but learnable skill. AI predictions become just another input into a complex, and potentially consequential, decision process. At the same time, there is much potential for dissensus given the complex relationships among those who collect, manage, and consume intelligence, not to mention the perennial risks of intelligence politicization.⁹⁹

A military organization needs to understand not only its adversaries but also itself. The prospects for AI are sometimes better for command and control (C2) than for ISR because friendly organizations and processes are easier to control, which produces more reliable reporting data. A lot of staff effort is consumed by searching for data, querying other organizations for data, and reanalyzing and reformatting data in response to emerging information requirements. AI can be used to integrate reporting data from disparate databases, helping to resolve contradictions and view activity in a “common operational picture.”¹⁰⁰ AI-produced decision aids can help personnel analyze unfolding battlefield conditions, run simulations of operational scenarios, and present options for military commanders to evaluate. For example, to deter-

97. James L. Regens, “Augmenting Human Cognition to Enhance Strategic, Operational, and Tactical Intelligence,” *Intelligence and National Security*, Vol. 34, No. 5 (2019), pp. 673–687, <https://doi.org/10.1080/02684527.2019.1579410>.

98. Minna Räsänen and James M. Nyce, “The Raw Is Cooked: Data in Intelligence Practice,” *Science, Technology, & Human Values*, Vol. 38, No. 5 (September 2013), pp. 655–677, <https://doi.org/10.1177/0162243913480049>.

99. Richard K. Betts, *Enemies of Intelligence: Knowledge and Power in American National Security* (New York: Columbia University Press, 2007); and Joshua Rovner, *Fixing the Facts: National Security and the Politics of Intelligence* (Ithaca, N.Y.: Cornell University Press, 2011).

100. Hoadley and Saylor, *Artificial Intelligence and National Security*, p. 12.

mine the best method for evacuating injured soldiers, militaries can use predictive models based on weather conditions, available routes, landing sites, and anticipated casualties.¹⁰¹ AI can be used to enhance computerized wargaming and combat simulations, offering more realistic models of “red team” behavior and more challenging training exercises.¹⁰² AI could potentially improve mission handoff between rotating units by analyzing unstructured text (i.e., passages of prose rather than standardized fields) in the departing unit’s reports.¹⁰³ Yet, as with intelligence, planning tasks cannot be simply delegated to machines. ISR and C2 reporting systems generate a mass of potentially relevant data, but they are hard to interpret, and associated metadata are often missing or misleading. In these situations, quality data may generate reliable predictions, but human intervention and interpretation is required throughout the decision process.

Human-machine teaming often entails not only task performance (i.e., balancing the cognitive load across people and AI) but also task design (i.e., adjusting the load as circumstances change). Viewed at a more granular level, a task that falls into the human-machine teaming category in our framework might be disaggregated into subtasks that fall into two of the framework’s other categories. That is, human practitioners will have to partition a complex decision task into either fully automated or fully human decision-making subtasks. This subdivision requires making mindful decisions about monitoring and controlling the risks of premature automation. For example, human-machine teaming in drone operations involves having both the drone and the drone operators perform certain tasks autonomously. The drone might automatically perform flying tasks (i.e., maintaining course and bearing or reacquiring a lost datalink), while human drone operators might deliberate over legal targeting criteria.

The overall partition (i.e., the location of the human in the loop) should be adjusted over time as conditions change, which will require humans to be mindful of how the division of labor between humans and machines relates to the task environment and the organizational mission. This balance will be further complicated by interdependencies across tasks and organizations, data access, interpretability, and interoperability issues, as well as competing priorities such as speed, safety, secrecy, efficiency, effectiveness, legality, cybersecurity, stability, adaptability, and so on. Importantly, as figure 1 shows, the

101. Benjamin Jensen and Ryan Kendall, “Waze for War: How the Army Can Integrate Artificial Intelligence,” *War on the Rocks*, September 2, 2016, <https://warontherocks.com/2016/09/waze-for-war-how-the-army-can-integrate-artificial-intelligence/>.

102. Kania, “Battlefield Singularity,” p. 28.

103. Spiegeleire, Maas, and Sweijs, *Artificial Intelligence and the Future of Defense*, p. 90.

organizational and political institutions that are exogenous to decision-making tasks establish the priorities for these different objectives. Humans are the ultimate source of judgment in all AI systems.

The Strategic Implications of Military AI

The central argument of this article is that machine learning is making prediction cheaper, which in turn makes data and judgment more valuable. This finding also means that quality data and clear judgment enhance AI performance. These conditions vary by decision task, but they are generally harder to meet in military situations given environmental and institutional complexities. Organizations that can meet them, however, may gain a competitive advantage. Human skills are central to this competitive advantage, and this has two important strategic implications.

First, military organizations that rely on AI have incentives to improve both data and judgment. These AI complements are sources of strength. At least one of them—judgment—relies wholly on human beings. Even when goals can be formally specified and pre-delegated for tasks in the automated decision-making category in our framework, humans must engineer the reward function, which they will likely revisit as they monitor system performance. AI adoption may radically change the distribution of judgment by altering who in an organization makes decisions and about what, but in all cases, humans are ultimately responsible for setting objectives, making trade-offs, and evaluating outcomes. There is little chance of this changing anytime soon given the technical state of the art. The other complement—data—also relies on human beings. Developing and implementing data policy necessitates negotiation between data producers and consumers. People also make nuanced judgments when architecting data infrastructure and managing data quality. AI systems can neither design themselves nor clean their own data, which leads us to conclude that increased reliance on AI will make human skills even more important in military organizations.

Second, and for the same reasons, adversaries have incentives to complicate both data and judgment. In a highly competitive environment, organizational strengths become attractive targets and potential vulnerabilities. Since predictable adversaries will play to AI strengths, intelligent adversaries will behave unpredictably. If AI creates military power in one area, adversaries will create military challenges in another. Facing an AI-empowered force, the enemy will attempt to change the game by either undermining the quality of predictions or making them irrelevant. Therefore, strategies to contest, manipulate, or disrupt data and judgment become more relevant as military competitors adopt

AI. The informational and organizational dimensions of war will continue to increase in salience and complexity. Again, this leads us to the conclusion that more military AI will make the human aspects of conflict more important.

This increased importance of human personnel challenges the emerging wisdom about AI and war. Many analyses either assume that AI will replace warriors for key military tasks or speculate that war will occur at machine speed, which in turn creates first-mover advantages that incentivize aggression and undermine deterrence.¹⁰⁴ The states that are first to substitute machines for warriors, moreover, are assumed to gain significant military advantages that will shift the balance of power toward early adopters. These outcomes are plausible, but they are based on problematic assumptions about AI substitutability. Conflicts based on AI complementarity may exhibit very different dynamics. We argue that it is more useful to consider the militarized contestation of AI complements (i.e., data and judgment) than to conceive of wars between automated military forces. Conflicts in which data and judgment are perennially at stake may be full of friction, controversy, and unintended consequences, and they may drag on in frustrating ways. In short, we expect the growing salience of data and judgment in war to subtly alter strategic incentives. As a result, AI-enabled conflicts are more likely to be decided by the slow erosion of resolve and institutional capacity than set-piece battles between robotic forces.

INFORMATION CONTESTS

The importance of information in war has been increasing for many decades.¹⁰⁵ The growth of ISR infrastructure—on and over the battlefield, at sea and underwater, and in orbit—has dramatically increased the volume and variety of data available to military organizations. Long-range precision weapons and high-bandwidth datalinks have also expanded the number of things that militaries can do with all these data, which in turn generates even more data about friendly operations. Yet, more and better data have not always translated into more effective military operations. The adoption of information technology throughout the past century has typically been accompanied by an increase in the complexity and geographical dispersion of military organizations. Data-intensive tasks that emphasize intellectual skills rather than physical fighting, such as intelligence, communications, and information operations,

104. For example, Payne, "Artificial Intelligence"; Horowitz, "When Speed Kills"; and Johnson, "Delegating Strategic Decision-Making to Machines."

105. Lindsay, *Information Technology and Military Power*, pp. 28–31; and Emily O. Goldman, ed., *Information and Revolutions in Military Affairs* (New York: Routledge, 2015).

have proliferated in military organizations. At the same time, advanced industrialized nations have asked their militaries to perform more complex operations. More complexity, in turn, increases the potential for disagreement and breakdown. Adversaries have also learned to offset the advantages of the ISR revolution by either adopting asymmetric tactics to blend in with civilian populations or exploiting the potential of space and cyberspace. As advances in battlefield sensors make it feasible to detect targets in near real time, enemy forces learn how to disperse, hide, and deceive.

In short, there may be more data in modern war, but data management has also become more challenging. Although U.S. weapons may be fast and precise, U.S. wars in recent decades have been protracted and ambiguous.¹⁰⁶ We argue that AI will most likely deepen rather than reverse these trends. Indeed, automation is both a response to and a contributing cause of the increasing complexity of military information practice.

Just as commanders are already preoccupied with C2 architecture, officers in AI-enabled militaries will seek to gain access to large amounts of data that are relevant to specific tasks in order to train and maintain AI systems. Units will have to make decisions about whether they should collect their own data organically or acquire shared data from other units, government agencies, or coalition partners. We expect many relevant databases to be classified and compartmented given the sensitivity of collection techniques or the content itself, which will complicate sharing. Units might also choose to leverage public data sources or purchase proprietary commercial data, both of which are problematic because nongovernmental actors may affect the quality of and access to data. As militaries tackle new problems, or new operational opportunities emerge, data requirements will change, and officers will have to routinely find and integrate new data sources. AI strategy will require militaries to establish data policies, and thus negotiating access to data will be an ongoing managerial—and human—challenge.

We contend that militaries will face not only data access but also data relevancy challenges. Heterogenous data-generating processes allow biases and anomalies to creep into databases. Although metadata may help to organize information processing, they are also vulnerable to data friction that only humans can fix.¹⁰⁷ Cleaning and curating data sources will therefore be as important as acquiring them in the first place. To the challenges of producing or procuring data must be added the challenges of protecting data. Just as supply

106. See, for example, Shimko, *The Iraq Wars and America's Military Revolution*.

107. Paul N. Edwards et al., "Science Friction: Data, Metadata, and Collaboration," *Social Studies of Science*, Vol. 41, No. 5 (October 2011), pp. 667–690, <https://doi.org/10.1177/0306312711413314>.

chains become attractive targets in mechanized warfare, data supplies will also become contested.

Overall, we expect the rise of AI to exacerbate the already formidable challenges of cybersecurity. Cybersecurity professionals aim to maintain the confidentiality, integrity, and availability of an organization's data. Two of these goals—integrity and availability—capture the AI requirements of unbiased and accessible data, as discussed above. The goal of confidentiality is also important insofar as data provide AI adopters with a competitive advantage. In commerce, AI companies often try to own (rather than buy) the key data that enable their machines to learn.¹⁰⁸ The military equivalent of this is classified information, which is hidden from the enemy to produce a decision advantage.¹⁰⁹ Military organizations will have strong incentives to protect the classified data that military AI systems use to learn. For the same reasons, adversaries will have incentives to steal, manipulate, and deny access to AI learning data. To date, most discussions of AI and cybersecurity have focused on a substitution theory of cybersecurity, that is, using AI systems to attack and defend networks.¹¹⁰ But we argue that a complementary theory of cybersecurity is just as if not more important. AI will require the entire military enterprise to invest more effort into protecting and exploiting data. If AI systems are trained with classified information, then adversaries will conduct more espionage. If AI enhances intelligence, then adversaries will invest in more counter-intelligence. If AI provides commanders with better information, then adversaries will produce more disinformation.

Inevitably, different parts of the bureaucracy will tussle among themselves and with coalition partners and nongovernmental actors to access and curate a huge amount of heterogeneous and often classified data. Organizations will also struggle with cyber and intelligence adversaries to maintain control of their own data while also conducting their own campaigns to collect or manipulate the enemy's data. To appreciate the strategic implications of AI, therefore, it is helpful to understand cyber conflict, most of which to date resembles espionage and covert action more than traditional military warfare. In fact, chronic and ambiguous intelligence contests are more common than fast and decisive cyberwar.¹¹¹ Military reliance on AI becomes yet another factor abet-

108. Agrawal, Gans, and Goldfarb, *Prediction Machines*, pp. 174–176.

109. Jennifer E. Sims, "Decision Advantage and the Nature of Intelligence Analysis," in Loch K. Johnson, ed., *The Oxford Handbook of National Security Intelligence* (New York: Oxford University Press, 2010).

110. For example, James Johnson, "The AI-Cyber Nexus: Implications for Military Escalation, Deterrence, and Strategic Stability," *Journal of Cyber Policy*, Vol. 4, No. 3 (2019), pp. 442–460, <https://doi.org/10.1080/23738871.2019.1701693>.

111. Joshua Rovner, "Cyber War as an Intelligence Contest," *War on the Rocks*, September 16, 2019,

ting the rise of cyber conflict in global affairs, and the (ambiguous, confusing, interminable, gray zone) dynamics of cyber conflict are likely to have a strong influence on the dynamics of AI conflict.

ORGANIZATIONAL COMPLEXITY

Just as AI militaries will struggle to procure, clean, curate, protect, and contest data, they will also struggle to inculcate, negotiate, and legitimate judgment. Indeed, the challenges of data and judgment go hand in hand. People will find it harder to interpret a flood of heterogeneous data. More complex data architectures will require managers to consider the trade-offs among competing objectives (i.e., confidentiality, integrity, and availability), which may invite bureaucratic controversy. Yet, judgment is even more fundamental for organizations that rely on AI because humans must both tell AI systems which predictions to make and determine what to do with the predictions once they are made. People who code valuations into autonomous systems will have enormous power because AI increases the scale of the impact of some human judgments. For example, individual car drivers make judgments about their own vehicle, whereas the encoded judgments for self-driving cars can affect millions of vehicles. Each instance of a given autonomous weapon system, similarly, will likely share algorithms and training data with others. When widely shared judgments are wrong, biased, or self-serving, then the AI systems guided by them can generate large-scale problems. Good judgment becomes particularly desirable as prediction gets better, faster, and cheaper.

A fundamental organizational challenge is to recruit, train, and retain the human talent required for human-machine teaming. We anticipate that AI systems will increase the influence of junior personnel, giving more leverage to their judgment and decisions. Yet, we also expect that the junior officers, non-commissioned officers, civilian employees, and government contractors who maintain and operate AI systems will struggle to understand the consequences of their actions in complex political situations. Gen. Charles Krulak highlights the role of “the strategic corporal” on twenty-first-century battlefields.¹¹² Krulak argues that operational complexity makes tactical actions more strategically consequential, for better or worse, which places a premium on the char-

<https://warontherocks.com/2019/09/cyber-war-as-an-intelligence-contest/>; Lennart Maschmeyer, “The Subversive Trilemma: Why Cyber Operations Fall Short of Expectations,” *International Security*, Vol. 46, No. 2 (Fall 2021), pp. 51–90, https://doi.org/10.1162/isec_a_00418; and Robert Chesney and Max Smeets, eds., *Cyber Conflict as an Intelligence Contest* (Washington, D.C.: Georgetown University Press, forthcoming).

112. Charles C. Krulak, “The Strategic Corporal: Leadership in the Three Block War,” *Marines Magazine*, January 1999, <https://apps.dtic.mil/sti/pdfs/ADA399413.pdf>.

acter and leadership ability of junior personnel. AI will further increase the burden of judgment on them. Forward personnel will have to see the predictions from AI systems, assess whether the data that created the predictions are reliable, and make value judgments about how and why automated systems can advance the mission. Furthermore, AI systems will require constant reconfiguration and repair as the context of human-machine teaming changes during actual operations. Military personnel have long engaged in field-expedient, bottom-up innovation.¹¹³ We expect personnel will likewise hack AI systems to improve mission performance, as they understand it, even as unauthorized modifications put them into conflict with system configuration managers elsewhere in the bureaucracy.¹¹⁴ It is important to emphasize the human capital requirements of combining a sophisticated understanding of the politico-military situation with the technical savvy to engineer AI in the field. The strategic corporal in the AI era must be not only a Clausewitzian genius but also a talented hacker. This may not be a realistic requirement.

The importance of human-machine teaming is increasingly appreciated in organizations that implement AI systems. Amid all the hype about AI and war, plenty of thoughtful work seeks to discern the relative advantages of humans and machines and to devise methods of pairing them together in order to improve decision-making.¹¹⁵ As the U.S. Department of Defense AI strategy states, “The women and men in the U.S. armed forces remain our enduring source of strength; we will use AI-enabled information, tools, and systems to empower, not replace, those who serve.”¹¹⁶ Yet, the strategy’s stated goal of “creating a common foundation of shared data, reusable tools, frameworks and standards, and cloud and edge services” is more of a description of the magnitude of the problem than a blueprint for a solution.¹¹⁷ As AI creates

113. Kollars, “War’s Horizon.”

114. On the general dynamics of military user innovation see Lindsay, *Information Technology and Military Power*, pp. 109–135.

115. Andrew Herr, “Will Humans Matter in the Wars of 2030?” *Joint Force Quarterly*, Vol. 77, No. 2 (2015), pp. 76–83, <https://ndupress.ndu.edu/Portals/68/Documents/jfq/jfq-77/jfq-77.pdf>; Mary L. Cummings, *Artificial Intelligence and the Future of Warfare* (London: Chatham House, Royal Institute of International Affairs, January 26, 2017); Development, Concepts and Doctrine Centre, “Human-Machine Teaming,” Joint Concept Note 1/18 (London: UK Ministry of Defence, May 2018), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/709359/20180517-concepts_uk_human_machine_teaming_jcn_1_18.pdf; and Mick Ryan, “Extending the Intellectual Edge with Artificial Intelligence,” *Australian Journal of Defence and Strategic Studies*, Vol. 1, No. 1 (2019), pp. 23–40, <https://www.defence.gov.au/ADC/publications/AJDSS/documents/volume1-issue1/Full.pdf>. NSCAI’s *Final Report*, 2021, also emphasizes human-machine teaming.

116. *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*, 2019, p. 4.

117. *Ibid.*, p. 7.

potential for large-scale efficiency improvements, it also creates potential for large-scale collective action problems. New military staff specialties are sure to emerge to manage data and judgment resources, creating new institutional equities and integration challenges. Perhaps even more challenging is the problem of nurturing trust among all the engineers, administrators, analysts, operators, and lawyers involved in designing, using, and repairing AI systems.¹¹⁸

As cheap prediction makes human judgment more vital in a wide variety of tasks, and as more judgment is needed to coordinate human-machine teaming, we anticipate that military bureaucracies will face complicated command decisions about why, and how, to conjoin humans and machines. Commercial firms that embrace AI often adjust their boundaries and business models by contracting out tasks involving data, prediction, and action (e.g., manufacturing, transportation, advertising, and service provision) while developing in-house judgment capacities that are too difficult to outsource.¹¹⁹ Military organizations, likewise, may find it advantageous to share specialized resources (sensors, shooters, intelligence products, and logistics) across a decentralized network of units, even as they struggle to make sense of it all. AI is thus part of a broader historical trend that has been described with terms like “network-centric warfare,” “joint force operations,” “integrated multi-domain operations,” and “interagency cross-functional teams.” The whole is more than the sum of its parts, but each part must exercise excellent judgment in how it leverages shared assets. Historical experience suggests that military interoperability and shared sensemaking are difficult, but not necessarily impossible, to achieve.¹²⁰ We thus expect military and political judgment will become even more difficult, diffused, and geographically distributed.

Indeed, the ongoing involvement of the “strategic corporal” in conversations about politico-military ends could end up politicizing the military. In the United States, as Risa Brooks argues, the normative separation of political ends from military means has some paradoxically adverse consequences: it enables service parochialism, undermines civilian oversight, and degrades strategic deliberation.¹²¹ Greater reliance on AI could exacerbate all these problems,

118. Roff and Danks, “Trust but Verify.”

119. Agrawal, Gans, and Goldfarb, *Prediction Machines*, pp. 170–178.

120. For example, C. Kenneth Allard, *Command, Control, and the Common Defense* (New Haven, Conn.: Yale University Press, 1990); and Scott A. Snook, *Friendly Fire: The Accidental Shootdown of U.S. Black Hawks over Northern Iraq* (Princeton, N.J.: Princeton University Press, 2000).

121. Risa Brooks, “Paradoxes of Professionalism: Rethinking Civil-Military Relations in the United States,” *International Security*, Vol. 44, No. 4 (Spring 2020), pp. 7–44, https://doi.org/10.1162/isec_a_00374.

precisely because AI is a force multiplier that requires military personnel to exercise greater judgment. Brooks's argument implies that an AI-intensive defense bureaucracy could become both more powerful and more politically savvy. If machines perform the bulk of data gathering, prediction, and tactical warfighting, then the judgments of human engineers, managers, and operators will be highly consequential, even as ethical questions of accountability become harder to answer. Some military personnel may be unable to perform at such a high level of excellence, as attested by the many scandals during the wars in Iraq and Afghanistan (from targeting errors to prisoner abuse). Increasing reliance on AI will magnify the importance of leadership throughout the chain of command, from civilian elites to enlisted service members.

If a military organization can figure out how to recruit, train, and retain highly talented personnel, and to thoroughly reorganize and decentralize its C2 institutions, such reforms may help to inculcate and coordinate judgment. Doing so would enable the military to make the most of human-machine teaming in war. If judgment is a source of military strength, however, then it may also be a political vulnerability. As organizational and political judgment becomes the preeminent source of strength for AI-enabled military forces, we expect that judgment will also become the most attractive target for adversaries. If AI relies on federated data and command structures, then adversaries will pursue wedge strategies to break up military coalitions.¹²² If the consensus about war aims depends on robust political support, adversaries will conduct disinformation and influence campaigns to generate controversy and undermine popular support.¹²³ If automated systems operate under tightly controlled rules of engagement, adversaries will attempt to manipulate normative frameworks that legitimize the use of force.¹²⁴ If AI enables more efficient targeting, the enemy will present more controversial and morally fraught targets to test political resolve.¹²⁵ As prediction machines make some aspects of military operations more certain, we argue that the entire military enterprise will become less certain.

122. See, generally, Timothy W. Crawford, "Preventing Enemy Coalitions: How Wedge Strategies Shape Power Politics," *International Security*, Vol. 35, No. 4 (Spring 2011), pp. 155–189, https://doi.org/10.1162/ISEC_a_00036.

123. See, generally, Thomas Rid, *Active Measures: The Secret History of Disinformation and Political Warfare* (New York: Farrar, Straus and Giroux, 2020).

124. Janina Dill, *Legitimate Targets? Social Construction, International Law, and US Bombing* (Cambridge: Cambridge University Press, 2014); and Ryder McKeown, "Legal Asymmetries in Asymmetric War," *Review of International Studies*, Vol. 41, No. 1 (January 2015), pp. 117–138, <https://doi.org/10.1017/S0260210514000096>.

125. Erik Gartzke and James Igoe Walsh, "The Drawbacks of Drones: The Effects of UAVs on Escalation and Instability in Pakistan," *Journal of Peace Research*, forthcoming.

Conclusion

It is premature to assume that AI will replace human beings in either war or any other competitive endeavor. To understand the impact of AI in any field, it is important to disaggregate decision-making into its components: data, judgment, prediction, and action. An economic perspective on AI views machine learning as more efficient prediction (and robotics as a more efficient action), which makes data and human judgment more valuable. This means that innovation in algorithms and computing power is necessary but not sufficient for AI performance. We have argued that the context of decision-making—where and how organizations use AI and for what purposes—determines whether automation is possible or desirable. The complementarity of data and judgment, in turn, has important implications for the preparation for and conduct of AI-enabled war.

We have argued that the strategic environment shapes the quality of data, and organizational institutions shape the difficulty of judgment, which gives rise to four different categories of AI performance in military tasks. Quality data and clear judgment enable “automated decision-making,” which is most feasible for bureaucratically constrained administration and logistics tasks. Low-quality data and difficult judgments, which are common in strategy and command tasks, necessitate “human decision-making.” Clear judgments applied to low-quality data create risks of “premature automation,” especially when AI systems are authorized to execute fire and maneuver tasks. Quality data and difficult judgments can be combined in “human-machine teaming,” which can be used to improve intelligence and planning tasks. We expect that many, if not most, practical military applications of AI are likely to fall into this last category. Even highly bureaucratized tasks that seem to fit in the “automated decision-making” category can require human judgment, especially when budget and personnel decisions are at stake or when resource scarcity creates difficult operational trade-offs. Likewise, highly nuanced command tasks that seem to fit in the “human decision-making” category can usually be broken down into a subset of tasks that might benefit from AI decision aids. Most practitioners who implement military AI systems are aware of the risks of “premature automation” in fire and maneuver, in part due to widespread apprehension about “killer robots.”¹²⁶ To determine the appropriate division of labor between humans and machines, therefore, humans must decide what

126. Roff, “The Strategic Robot Problem.”

to predict, and they must create data policies and AI learning plans that detail who should do what with such predictions.¹²⁷ The dynamic circumstances of military operations will require ongoing finessing of the human-machine teaming relationship.

Although we agree with the conventional wisdom that AI is potentially transformative, we disagree about what that transformation might be.¹²⁸ In general, we expect that the strategic, organizational, and ethical complexity of warfare will increase in the AI era. When cheaper prediction is applied in a political context that is as challenging and uncertain as warfare, then quality data and sound judgment become extremely valuable. Adversaries, in turn, will take steps to undermine the quality of data and judgment by manipulating information and violating expectations. Correcting for adversarial countermeasures will further increase the complexity of judgment, which exacerbates the inherent friction and frustration of war.

We must reemphasize that our focus throughout has been on narrow AI, particularly the improvements in machine learning that have led to better, faster, and cheaper predictions. We contend that the recent advances in AI that have led to media attention, commercial applications, and anxiety about civil liberties have very little to do with AGI. Some experts believe that AGI will eventually happen, but this is not what all the current AI hype is about.¹²⁹ Other experts like Brian Cantwell Smith are outright pessimistic: “Neither deep learning, nor other forms of second-wave AI, nor any proposals yet advanced for third-wave, will lead to genuine intelligence.”¹³⁰ Indeed, the “intelligence” metaphor is very misleading when it comes to understanding what machine learning actually does.¹³¹ Advances in narrow AI, by contrast, have led to better, faster, and cheaper predictions. Such AI systems are task-specific.

127. Heuristics are provided in Agrawal, Gans, and Goldfarb, *Prediction Machines*, pp. 123–151.

128. See, for example, Horowitz, “Artificial Intelligence, International Competition, and the Balance of Power”; and Payne, “Artificial Intelligence.”

129. Daniel Kahneman, “Comment on ‘Artificial Intelligence and Behavioral Economics,’” in Agrawal, Gans, and Goldfarb, eds., *The Economics of Artificial Intelligence*, pp. 608–610. Gary Marcus estimated that AGI would arrive between thirty and seventy years from now. See Shivon Zilis et al., “Lighting Round on General Intelligence,” panel presentation at Machine Learning and the Market for Intelligence Conference, Creative Destruction Lab, University of Toronto, October 26, 2017, YouTube video, 13:16, https://www.youtube.com/watch?v=RxLIQ_j_BMhk.

130. Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge: Massachusetts Institute of Technology Press, 2019), p. xiii. See also Harry Collins, *Artificial Experts: Social Knowledge and Intelligent Machines* (Cambridge: Massachusetts Institute of Technology Press, 1990); and Meredith Broussard, *Artificial Unintelligence: How Computers Misunderstand the World* (Cambridge: Massachusetts Institute of Technology Press, 2018).

131. A less anthropocentric definition requires a longer discussion about the meanings of intelligence, autonomy, and automation. See Heather M. Roff, “Artificial Intelligence: Power to the People,” *Ethics & International Affairs*, Vol. 33, No. 2 (2019), pp. 124–140, <https://doi.org/10.1017/S0892679419000121>.

If AGI becomes a reality, then such a machine would also provide its own judgment. AGI would be able to perform the entire decision cycle by itself. In that case, it is not at all clear what role humans would have in warfare beyond suffering the consequences of war.¹³² We argue that AGI speculation carries the theme of AI substitution to an extreme, whereby a machine would be able to outwit, overpower, and eliminate any actor who tried to prevent it from accomplishing its goal.¹³³ This doomsday scenario is often likened to the “Sorcerer’s Apprentice” segment from the movie *Fantasia*, in which the eponymous apprentice, played by Mickey Mouse, enchants a broom and directs it to fetch water from the well. As Mickey falls asleep, the broom ends up flooding the entire castle. Mickey awakes with alarm and desperately tries to chop up the broom, but this only results in more and better brooms that overwhelm his abilities. An eminently useful tactical task turns into a strategic disaster because of a poorly specified objective. Opinions vary on whether the superintelligence threat should be taken seriously.¹³⁴ Nevertheless, the Sorcerer’s Apprentice scenario dramatizes the importance of judgment for any type of AI. An AI that only cares about optimizing a goal—even though that goal was defined by a human—will not consider the important pragmatic context that humans may care about.

We have defined judgment narrowly in economic terms as the specification of the utility function. The rich concept of judgment, however, deserves further analysis. Just as decision-making can be disaggregated into its components, judgment might also be disaggregated into the intellectual, emotional, and moral capacities that people need to determine what matters and why. Military judgment encompasses not only the Clausewitzian traits of courage, determination, and coup d’oeil, but also a capacity for fairness, empathy, and other elusive qualities. Some wartime situations merit ruthlessness, deviousness, and enmity, while others call for mercy, candor, and compassion. To these character traits must be added the engineering virtues of curiosity, creativity, and elegance insofar as personnel will have to reconfigure AI systems in the field. We expect that the general logic of complementarity will still apply at this more fine-grained level. Any future AI that is able to automate some aspects of judgment, therefore, will make other aspects even more valu-

132. For speculation on the consequences of an AGI that is able to formulate and execute politico-military strategy, see Kenneth Payne, *Strategy, Evolution, and War: From Apes to Artificial Intelligence* (Washington, D.C.: Georgetown University Press, 2018).

133. Bostrom, *Superintelligence*; and Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* (New York: Viking, 2019).

134. For discussion see Nathan Alexander Sears, “International Politics in the Age of Existential Threats,” *Journal of Global Security Studies*, Vol. 6, No. 3 (September 2021), <https://doi.org/10.1093/jogss/ogaa027>.

able. Furthermore, the rich phenomenology of judgment, which AI makes more valuable, has important implications for professional military education. More technology should not mean more technocracy. On the contrary, personnel would be wise to engage more with the humanities and reflect on human virtues as militaries become more dependent on AI. In general, reliance on AI will tend to amplify the importance of human leadership and the moral aspects of war.

In the end, we expect that more intensive human-machine teaming will result in judgment becoming more widely distributed in military organizations, while strategic competition will become more politically fraught. Whatever the future of automated warfare holds, humans will be a vital part of it.