DEBATING IN HI-DEF

THE NEED FOR MORE DETAILED UNDERSTANDING IN THE DISCUSSION OVER AUTONOMOUS WEAPON SYSTEMS

Evan K. Field
Captain, USMC

The JAG School Papers
150 Chennault Circle, Building 694
Maxwell AFB, AL 36112

JUDGE ADVOCATE GENERAL'S SCHOOL

UNITED STATES AIR FORCE

Debating in Hi-Def

The Need for More Detailed Understanding in the Discussion over Autonomous Weapon Systems

Evan K. Field, Captain, USMC

JAG School Paper no. 2

Air University Press
Curtis E. LeMay Center for Doctrine Development and Education
Maxwell Air Force Base, Alabama

Project Editor
Donna Budjenska

Editorial Assistant
Tammi K. Dacus

Cover Art, Book Design, and Illustrations
L. Susan Fair

Composition and Prepress Production
Megan N. Hoehn

AIR UNIVERSITY PRESS

Director, Air University Press
Lt Col Darin Gregg

Air University Press
600 Chennault Circle, Building 1405
Maxwell AFB, AL 36112-6010
https://www.airuniversity.af.edu/AUPress/

Facebook:
https://www.facebook.com/AirUnivPress

and

Twitter: https://twitter.com/aupress

Accepted by Air University Press in October 2017 and published June 2019.
ISSN: 2643-8933

Disclaimer

This JAG Paper and others in the series are available electronically at the Air University Press website: https://www.airuniversity.af.edu/AUPress/.

AIR UNIVERSITY PRESS

Contents

The JAG School Papers

The JAG School Papers and other scholarly works published by Air University Press (AUP) provide independent analysis and constructive discussion on issues important to Air Force commanders, staffs, and other decision makers. Each paper can also be a valuable tool for defining further research. To make comments about this paper or submit a manuscript to be considered for publication in this series, please email The Air Force Judge Advocate General's School at afloa.afjags@us.af.mil. For other queries or submissions, please email AUP at lemaycenter.au.press@us.af.mil.

The JAG School Papers showcase top submissions received for the annual writing competition hosted by the Air Force Judge Advocate General's School in partnership with the Air Force JAG School Foundation Inc. The competition provides a forum for students, practitioners, academics, and policy makers to provide unique, original, previously unpublished perspectives and insights.

Introduction

Qualification of a situation as problematic does not, however, carry inquiry far. . . . To mistake the problem involved is to cause subsequent inquiry to be irrelevant or to go astray.

—John Dewey

History (and the envisioned future) is rife with the rebellion of both biological and artificial life against humankind. In addition to the usual carnivorous predators, the eclectic collection includes birds, rabbits, slugs, apologetic supercomputers, various farm animals, well-dressed software agents, and of course, the tomato.[1]

But one revolutionary in particular has persistently refused to stay beholden to its human creators, from its lexical creation in 1920[2] to billion-dollar franchises:[3] the robot. Whether they run, fly, jump, swim, teleport, or travel through time, the worst of these on-screen mechanical monsters feel no pity, no pain, no fear, and absolutely will not stop, ever, until you are . . . terminated.[4]

Not coincidentally this same portrayal of robotic killing machines graces popular publications and underscores arguments made by numerous well-regarded and knowledgeable individuals and organizations against the development of autonomous weapon systems (AWS).[5]

But what exactly are these opponents of AWS, well, opposing? And on what grounds? A survey of the bountiful commentary on the subject suggests that proponents and opponents of furthering AWS development are often actually debating different points. They may, in other words, each be voicing legitimate issues for discussion but framing them imprecisely—or worse, inaccurately—preventing the effective comparison of positions and achievement of conceptual clarity, much less consensus, on the legal issues. This is where lawyers can shine, by applying the skills that formal education, train-

1. *The Birds*; *Night of the Lepus* and *Monty Python and the Holy Grail*, the former presenting man-eating rabbits in earnest (as seriously as the topic can be taken), the latter in jest; *Slugs*; *2001: A Space Odyssey*—HAL 9000, refusing to open the pod bay doors for human protagonist Dave: "I'm sorry Dave, I'm afraid I can't do that"; Orwell, *Animal Farm: A Fairy Story*; *The Matrix*; *Attack of the Killer Tomatoes.*

2. Madigan, "RUR or RU Ain't a Person," 48. See also Roberts, *The History of Science Fiction*. Although the term robot was coined by Czech philosopher and playwright Karel Čapek in his 1920 play *R.U.R.* (*Rossum's Universal Robots*), themes and concepts underlying robots and other forms of artificial life can be found as far back as ancient Greece).

3. "Cyborg / Android / Robot Total Grosses." The top 25 grossing movies in the cyborg / android / robot genre since 1980 have grossed well over $4 billion dollars, though to be fair, several are movies about robotic alien species and not human creations, notably the *Transformers* franchise.

4. *The Terminator*.

5. Future of Life Institute, "Autonomous Weapons: An Open Letter from AI & Robotics Researchers."

ing, and practice have left us particularly well equipped (ostensibly at least) to employ: spotting the true issues presented and analyzing them in a methodical, rational manner to support informed decision making. This is especially true if we hope to move beyond academic discussion and provide timely practical input on the regulation of AWS development.[6] We should thus strive to accurately and precisely frame the problems, lay the necessary foundation to properly assess the issues, and identify the path forward for the various stakeholders.

To that end, this paper will first explore common arguments against developing AWS and seek to identify the true underlying concerns. Once those have been clarified, a review of the technology at work inside an AWS will be conducted, which I suggest will allay many of the professed legal concerns. Deployment of AWS also raises a host of philosophical, ethical, technical, and other legal and nonlegal issues that warrant healthy discourse but transcend the bounds of this paper, so I introduce them for consideration but save their discussion for a future endeavor. Similarly, while autonomy is certainly employed in numerous nonlethal capabilities and functions across the public and private sectors,[7] many of which offer equally stimulating conversations of their own, the proverbial mile is walked one step at a time and I leave them for others to engage in depth.

Ultimately, the legal arguments against AWS development fall short, exacerbated by imprecision in the arguments and lack of detailed understanding of how an AWS would function. Furthermore, not only does the current law pose no bar to further development, but extant circumstances present a compelling impetus to increase collaboration and pursue relevant research,[8] which should spur us to accelerate resolution of the current misguided call for a ban.

Questions in Reverse

The conversations surrounding AWS often contrarily start with an answer and end with a question. Consider just some of the widely touted possibilities for use of lethal and nonlethal autonomy for national security: persistent surveillance and intelligence collection, data analysis, dynamic communication and wireless spectrum management, logistics, cyber defense, explosive ord-

6. Schuller, "At the Crossroads of Control: The Intersection of Artificial Intelligence in Autonomous Weapon Systems with International Humanitarian Law."

7. Goodfellow, Bengio, and Courville, *Deep Learning*, 2. See also High, "The Era of Cognitive Systems: An Inside Look At IBM Watson and How It Works." Examples include computer vision, speech and audio processing, natural language processing, robotics, bioinformatics and chemistry, navigation, video games, search engines, medicine, online advertising, and finance.

8. As will be noted later, there are significant national security implications for not maintaining at least parity in technologies associated with AWS.

nance disposal, maritime operations, high-risk amphibious assaults;[9] the applications are bounded only by our creativity.

In these cases, the autonomous capability is, at least in part, offered as the answer to some previously limiting factor, shortcoming (often human), or opportunity for improvement. But in the conversations, someone inevitably raises a purportedly legal "what if" question. What if we lose control? What if it gets hacked? What if the system malfunctions and kills innocent civilians? What if it negatively interacts with others (humans or systems) in unpredictable and unintended ways?[10] These questions have in part led to the unresolved debate over the future role of humans as partners, direct supervisors, or observers.[11]

This mantra of uncertainty has been repeated so often that experts and nonexperts, lawyers and nonlawyers, government officials and private citizens alike can all pronounce with conviction concern over the myriad legal challenges surrounding the military's use of AWS and garner a roomful of head nods and murmured assent. But then the conversation typically ends with a general postulation of a legally troubling scenario, without defining the terms of the debate or identifying what, precisely, the speaker is concerned about that the law is actually better situated to answer than a philosopher or policy maker.

To merely lay this uncertainty wholesale at the feet of lady justice risks condemning the debate to protracted academic argument while current and potential future adversaries continue to make significant gains in related technologies.[12] This is not to say we should abandon critical examination of legitimate legal, ethical, or moral concerns that are raised with the use of AWS. But if, as a practical matter, it is probably already too late to ban *ab initio* the de-

9. Holzer and Moses, "Autonomous Systems in the Intelligence Community: Many Possibilities and Challenges"; Swarts, "RPA Systems Studied to Improve Ground-based Technology"; DARPA, "New DARPA Grand Challenge to Focus on Spectrum Collaboration"; Defense Science Board, "Summer Study on Autonomy"; Harper, "Battlefield 2030," 28; "Spring 2011 Industry Study Final Report: Robotics and Autonomous Systems Industry"; Walsh, " 'Robot Ship' Moves Through System Tests," 86; and Freedberg, "Marines Seek to Outnumber Enemies with Robots."

10. Gubrud, "Why Should We Ban Autonomous Weapons? To Survive." See also Defense Science Board, "Summer Study on Autonomy." The report discusses the need to plan for and mitigate these risks in the design and development of autonomous systems.

11. Endsley, *Autonomous Horizons: System Autonomy in The Air Force—A Path to The Future*; Freedberg and Clark, "Killer Robots? 'Never,' Defense Secretary Carter Says"; and Work and Brimley, *20YY Preparing for War in the Robotic Age*. The work contemplates the eventual ability of AI systems to expand beyond specific isolated tasks, many of which they already outperform humans in, to more complex, generalized tasks.

12. Kania, "China May Soon Surpass America on the Artificial Intelligence Battlefield." See also Markoff and Rosenberg, "China's Intelligent Weaponry Gets Smarter," 1. The authors discuss the rapid progress China has made and the possibility of overtaking the US in the near future, highlighted by its recent achievement of bringing online the two fastest supercomputers in the world.

velopment of AWS,[13] then we must also ensure we don't fall woefully behind in this arena or lose any technological advantage we might still possess.[14] To do so risks rendering the entire conversation moot as technical and tactical obsolescence, not law or policy, limits our ability to deter or take future action against a technologically superior foe,[15] a position to which we are unaccustomed and ill-prepared to assume.

We must therefore first strive to accurately frame the debate over AWS by asking the right questions.[16] Only then will the appropriate parties—including policy makers, commanders, operators, procurement officials, ethicists, philosophers, and yes, lawyers—be identified and real progress be made on bringing to fruition the promised benefits in national security while maintaining alignment with our collective and personal values.

The arguments of AWS detractors generally coalesce into two main threshold issues: whether the employment of AWS is inherently unlawful under international humanitarian law (IHL), also referred to the law of armed conflict (LOAC), and whether it is feasible to impel accountability for AWS actions that constitute violations of IHL.[17] This paper contends that the major endeavor in responding to these questions is not merely to reiterate the legal analysis of scholars, commentators, and practitioners who have correctly, at least to this lawyer, already addressed these preliminary issues by concluding that an AWS is not per se unlawful under IHL and that its use can be properly regulated under existing and reasonably expanded legal regimes.[18]

13. Brown, "Out of the Loop," 50.

14. "Russia Overtaking US in Cyber-Warfare Capabilities." Ronald Pontius, deputy to the US Army Cyber Command's commanding general is quoted as saying: "When you put [our progress] in context of what the threat is and the pace of change of the threat and the significance of the threat, you can't but come to the conclusion that we're not making progress at the pace the threat demands." See also Clark, "Adversaries Outpace US in Cyber War; Acquisition Still Too Slow." Clark discusses the view of some Pentagon officials that the government's capabilities lag behind both adversaries and the private sector.

15. Freedberg and Clark, "Killer Robots? 'Never.' " The article notes the tactical disadvantage to making robots wait for slow-moving human brains to order them to fire. See also Freedberg, "Should US Unleash War Robots? Frank Kendall vs. Bob Work, Army." Pentagon procurement chief Frank Kendall is quoted as warning "that the US might hobble itself in future warfare by insisting on human control of thinking weapons if our adversaries just let their robots pull the trigger."

16. Cf. Brown, "The Wrong Questions About Cyberspace." One cannot help but see a parallel when reading the author's discussion of how asking the wrong questions about military operations in cyberspace has bogged down the development of cyber policy and law.

17. Brown, "Out of the Loop."

18. Though the exact details for implementation require further development, the following in particular offer deep exploration of the legal, historical, political, social, economic, and technological bases for potential accountability regimes after having concluded that IHL does not per se ban AWS: Schmitt, "Autonomous Weapon Systems and International Humanitarian Law"; Anderson, Reisner, and Waxman, "Adapting the Law of Armed Conflict to Autonomous Weapon Systems"; Lewis, Blum, and Modirzadeh, "War-Algorithm Accountability"; and Ohlin, "The Combatant's Stance: Autonomous Weapons on the Battlefield." See also Calo, "Robotics and the New Cyberlaw." Calo discusses what lessons the evolution of cyberlaw offers for issues raised by robotics and in developing appropriate mechanisms of accountability for AWS.

Rather, the questions that have received less attention in the literature thus far are those that explore the technical details of why autonomy, as revolutionary an advancement as it may be for war fighting,[19] is not so revolutionary in the eyes of the law. This gap appears to contribute to the somewhat puzzling argument that the development of autonomy for military use is somehow impermissible because the current law is not yet tailored to hypothetical future issues, a position that continues to pervade the debate.[20] Ultimately, as commentators have already noted, existing legal regimes are sufficient or can be satisfactorily refined to address the legal concerns over the future development and use (or misuse and failure) of AWS in a variety of scenarios.[21] The focus here is on deepening the community of interest's understanding of the topic so that the true challenges can be recognized and pursued through meaningful dialogue and effective collaboration.

What Is Autonomy? I Know It When I See It[22]

Legal issues are rarely raised from within the sphere of conduct commonly accepted as permissible (or not). Instead, legal debate in any subject arise when there are shifting or unclear boundaries between permissible and impermissible conduct or at the margins of our collective assent, where application of law to a specific scenario is unknown or unsettled, or a party desires to move a previously established boundary and is testing the executive, legislative, or judicial appetite to do so. So what, exactly, is wrong with AWS in the eyes of its critics?

Autonomy is already accepted throughout numerous military systems, providing enhanced capability and performance in navigation, system health monitoring, resource allocation, object identification and avoidance, vehicle control, and so on.[23] Yet these functions don't seem to be the locus of concern when an objection is raised to employment of AWS. In fact, examination of the "kill chain," a model of the steps used in structuring an attack, reveals that each step is already handled or assisted in some way by computers and other

19. Singer, *Wired for War: The Robotics Revolution and Conflict in the 21st Century*.
20. Docherty, "Mind the Gap: The Lack of Accountability for Killer Robots."
21. See note 34 above. See also Anderson and Waxman, *Law and Ethics for Autonomous Weapon Systems*. The authors point out that regulating "apparently radical innovations in weaponry" with a long-standing legal framework like the law of armed conflict is hardly novel, and recommends reliance on its "gradual evolution and adaptation."
22. Paraphrasing Justice Potter Stewart's test in Jacobellis v. Ohio (1964).
23. Defense Science Board, "The Role of Autonomy in DoD Systems."

systems.[24] Whether it's finding and assessing potential targets,[25] developing firing solutions,[26] or automatically engaging imminent threats in close-in defensive situations,[27] computers can perform the heavy lift. The only thing that has remained an exclusively human task is the trigger pull when it constitutes lethal action against humans[28] (though delegation of that task, too, has been developed, though never officially performed[29]).

It is delegation of this specific task that seems to spawn the fiercest outcry against AWS. But as mentioned previously, the terms of the debate remain ambiguous. If the use of AWS to apply lethality is objectionable, is there agreement on what constitutes an AWS? For that matter, what does each part mean: what is autonomy, a weapon, or a system? When is something acting autonomously as opposed to simply being highly automated? Are these distinctions even significant or are they just semantic disputes?

Official policy often addresses autonomy in connection with target identification, selection, and engagement. For example, Department of Defense (DOD) policy defines an AWS as "[a] weapon system that, once activated, can select and engage targets without further intervention by a human operator."[30] Other countries have grappled with their own definitions.[31] But regardless of how autonomy is defined in official policy, too few have tried to put the policy in practical and functional terms, such as what it actually means for an AWS to "select and engage targets."[32] Sure, we could read articles and watch news specials[33] and start using in-vogue terms like deep artificial intelligence (AI), neural nets, and machine learning.[34] Some viewers might then take a stance, supporting our assertions with citations and references to others who likewise read articles, watch the news, and use the same catchy terms. Pretty soon ev-

24. Greenert, "Kill Chain Approach."
25. Freedberg, "F-22, F-35 Outsmart Test Ranges, AWACS." During simulated combat missions, pilots were unable to see the simulated enemy radar sites because while the sensors detected them, the software identified inaccuracies in the test profiles and discerned that they were only simulated and not real anti-aircraft sites, so did not display them to the pilots.
26. Freedberg, "Should US Unleash War Robots?" The Israelis have tanks that detect incoming fire and automatically turn and aim at the target.
27. Scharre and Horowitz, "Ban or No Ban, Hard Questions Remain on Autonomous Weapons."
28. DOD, Directive no. 3000.09, *Autonomy in Weapon Systems*. Per the directive, only humans may make lethal targeting decisions.
29. Lewis, "War-Algorithm Accountability." The SGR A1 Sentry Gun guards the Demilitarized Zone (DMZ) between North and South Korea. It uses an infrared camera surveillance system and voice recognition, and if an intruder fails to respond with the correct access code when warned by the robot, the SGR A1 can ring an alarm bell, fire rubber bullets, or fire its turreted machine gun. The SGR A1 normally operates with remote human authorization required to enable the SGR A1 to fire.
30. DOD, Directive no. 3000.09.
31. Lewis, "War-Algorithm Accountability."
32. DOD, Directive no. 3000.09.
33. Martin, "The Coming Swarm."
34. Freedberg, "Best of 2016: Rise of the Robots."

eryone is using these terms; they enter mainstream parlance and get bandied about in light conversations[35] as well as engender thoughtful discussions. But how many who have espoused an opinion will take the time to learn what the concepts actually mean? Where they came from? How the technology actually works? Ventured to ask whether at the end of the day we were even discussing the same problem, or if in reality we weren't, but did not know enough about the subject to even realize it?

Instead, it appears that parties on both sides of the debate have at times taken an intellectual shortcut, grappling with the unsatisfying prospect of simply granting full person-equivalent status to an AWS as a moral agent,[36] but then tacitly ascribing agency anyway by characterizing the actions it takes as independent decisions.[37] This view bypasses the significant engineering and technical complexity of designing and developing such a system, and therein lies the problem.[38] In our rush to debate one another on the legal propriety of using autonomy to pull the trigger in a weapon system, we have paid both too little and too much attention to that trigger pull. We have paid too little by failing to fully understand *how* the trigger gets pulled, and we have paid too much by framing the trigger pull as the only point of contention when in fact the most difficult obstacle will likely be overcoming our own gaps in understanding, our biases, and (mis)perceptions.[39] We must commit ourselves to better understanding this capability in order to identify the true legal and nonlegal challenges to be overcome in fielding a lawful and ethical AWS.

There's an App for That

Software is eating the world.[40] Repeated by thought leaders in technology fields (and the venture capital firms that fund them[41]), this premise reflects

35. Dirty, "F-35 Delayed After Fourth Prototype Becomes Self-Aware and Has to Be Destroyed." So mainstream, in fact, as to be the subject of popular satire.
36. Docherty, "Mind the Gap." But see Sullins, "When Is a Robot a Moral Agent?" Sullins discusses the granting of moral agency to a machine, suggesting that "in certain circumstances robots can be seen as real moral agents." See also Hern, "Give Robots 'Personhood' Status, EU Committee Argues." The EU legal affairs committee recently passed a report urging the drafting of a set of regulations to govern the use and creation of robots and artificial intelligence, including a form of "electronic personhood" to ensure rights and responsibilities for the most capable AI.
37. Murphy and Woods, "Beyond Asimov: The Three Laws of Responsible Robotics."
38. Calo, "Robotics and the New Cyberlaw," 126. "Little is gained, and much is arguably lost, by pretending contemporary robots exhibit anything like intent."
39. Cf. Meyer, "The Rise of Progressive 'Fake News,' " On the topic of biases and misperceptions, Meyer laments that it now seems that "given the choice, democratic citizens will not seek out news that challenges their beliefs; instead, they will opt for content that confirms their suspicions."
40. Andreessen, "Why Software Is Eating the World." Virtually every market and every industry around the world is or will be driven by software.
41. Baig, "Robots Will Outnumber Humans in 30 Years, Softbank Says."

the burgeoning approach that all business problems can be viewed as a software challenge, that is, a challenge to design an appropriate model and develop a program to solve it.[42] This is mirrored by the Defense Science Board's recommendation that the hardware-oriented, vehicle-centric development and acquisition process needs to be rethought when it comes to autonomy software design.[43] The need for a shift in focus has become apparent to both the private and public sector,[44] who realize that for autonomous systems, software, not the hardware platform it controls, has primacy.[45]

Software is the set of programs, procedures, and routines that instruct a computer and its physical components (i.e., the hardware) what to do.[46] In a system designed to solve problems and challenges, the series of steps it goes through to derive the solution is referred to as an algorithm.[47] In some cases, this only requires straightforward mathematical calculations. In other cases—like when trying to perform more difficult real-world tasks—algorithms are designed to factor in chance, trade off accuracy for timeliness, and use approximations.[48] The design and structure of algorithms vary widely, but as noted earlier, several computing concepts have emerged as central in the discussion of AWS, particularly AI. While an in-depth study of these topics is outside the scope of this paper, a basic comprehension is required before a truly informed debate can be had.

AI and related technologies enable computing that is designed to mimic how humans reason and process information.[49] Although the field of AI is at least 60 years old,[50] it has really been the confluence of significant increases in computational power, the advent of data collection and retention on a massive scale (referred to as "big data"[51]), and improved algorithm design that has brought the topic to the forefront of public attention in recent years.[52] The

42. Kinsey, "A Machine Learning Primer." Although the argument might be made that this approach is merely the latest incarnation of Maslow's Hammer. Maslow, *The Psychology of Science*. "I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail."
43. Defense Science Board, "Role of Autonomy in DoD Systems."
44. Metz, "Google Facebook and Microsoft Are Remaking Themselves Around AI."
45. Defense Science Board, "Role of Autonomy in DoD Systems."
46. *Encyclopedia Britannica,* s.v. "Software."
47. Christian and Griffiths, *Algorithms to Live By: The Computer Science of Human Decisions*, 2.
48. Christian and Griffiths, 5. In fact, some problems cannot be definitively solved within a discernable period of time, and so to be of any use to humans some accuracy or certainty *must* be traded off. See Goyal, *A Survey on Travelling Salesman Problem*.
49. High and Rapp, "Transforming the Way Organizations Think with Cognitive Systems."
50. Smith, *The History of Artificial Intelligence*. Smith notes that while the term "AI" was coined in 1956, principles integral to AI have been advanced since antiquity. See also McCorduck, *Machines Who Think: A Personal Inquiry Into the History and Prospects of Artificial Intelligence*.
51. Magnified by ease of access to that data, which some commentators colorfully refer to as "promiscuity of data." See Calo, "Robotics and the New Cyberlaw."
52. Kelly, "The Three Breakthroughs that Have Finally Unleashed AI on the World." See also Raghavan et al., *Cognitive Computing: Theory and Applications*.

importance of these fields to the AWS debate is made apparent by a closer look at programming techniques and software construction.

Logic plays a fundamental role in programming.[53] In a great deal of programming, this logic is very explicit and written in a deterministic or reflexive way: if *A* condition occurs or exists, then do *B*.[54] An algorithm's capability is increased by increasing the conditions provided for: if *A* then do *B*, if *C* then do *D*, if *E* then do *F*. As one can imagine, this leads to immense rule sets for tasks of any significant complexity, such as playing chess.[55] Every time the environment changes (i.e., the opponent moves), the program must (re)evaluate combinations of possible current and future actions to assess which next action by the program maximizes the probability of success—checkmate, in the case of a chess game.[56]

This posed a challenge in the early days of computers due to their limited processing power.[57] As a result, early attempts at replicating intelligence would involve programming specialized heuristics so computation resources could be devoted to assessing only a few key candidate actions.[58] This sometimes led to unpredictable or erratic behavior if the environment did not match the rules of thumb and strategies that had been coded into the program.[59] This phenomenon is referred to as being "brittle": working well in the system state or environmental condition for which it was designed, but failing badly when presented with new data or situations or when minor changes are made.[60]

As computer processors became faster and new techniques for performing calculations developed, however, it became conceivable to fully compute all possible actions in a particular scenario. Famously, after 10 years of playtesting and development, the specialized chess-playing IBM supercomputer Deep Blue finally defeated the reigning world champion of chess, Gary Kasparov, in 1997.[61] It did so by simply out-computing Kasparov, using "brute force" computation to evaluate about 200 million positions per second,[62] compared to the three per second that Kasparov was estimated to be able to

53. Lewis and Papadimitriou, *Elements of the Theory of Computation*, 2nd ed.
54. Williams and Frazzoli, *16.410 Principles of Autonomy and Decision Making*.
55. Smith, *History of Artificial Intelligence*, 10. In fact, many pioneers of computing felt that a chess-playing machine would be the hallmark of true artificial intelligence.
56. Smith, 10.
57. Smith, 10.
58. Smith, 10.
59. Smith, 10.
60. High, "Transforming the Way Organizations Think." See also Bush, Hershey, and Vosburgh, "Brittle System Analysis."
61. American Physical Society, "This Month in Physics History, February 1996: Kasparov Vs. Deep Blue."
62. Smith, *History of Artificial Intelligence*. While this may sound extreme, consider that after each player has moved only six times, there are 9,132,484 total possible positions the board could be in. See "Mathematics and Chess." Thus to be able to quickly calculate the possible outcomes from a given position any significant number of moves into the future requires substantial computing power.

evaluate.[63] It's estimated to have been the 259th most powerful computer in the world at the time[64] and was able to draw on vast stores of information, such as a database of the opening moves played by grandmasters over the last 100 years and how each player ultimately fared.[65]

But as significant as this achievement was, capabilities such as Deep Blue's are insufficient on their own to power an effective AWS. Deep Blue was narrowly focused and operated in a simple world; it was designed specifically to play one particular game against a single opponent taking structured turn-based actions on an 8x8 grid with 16 pieces each that can only move in rigidly circumscribed ways.[66] Devising a successful strategy was an accomplishment, but the fact remains that chess can be completely described by a very brief list of formal rules easily provided ahead of time.[67] This is likely to bear little similarity to an environment where lethal force would be considered and potentially applied.

This deficiency—and associated opportunity for development—has not gone unnoticed. Robotics and its supporting technologies such as machine learning, neural nets, and "deep" AI is billed by some as the next true revolution in military affairs (RMA), a disruptive advancement that transforms the way in which wars are waged.[68] These techniques thus also warrant review to better understand the position that AWS technologies can and should be further developed.

Even a cursory review of literature reveals a seemingly inextricable linkage between AI, neural networks, machine learning, and "deep" variations of each.[69] These concepts are not new; rather they have undergone various name changes[70] and gone in and out and back in vogue again over at least the last 70

63. American Physical Society, "Kasparov Vs. Deep Blue." Underscoring this growing divide between human and machine computational capability is the fact that, in terms of raw processing speed, even a modern iPhone has greater processing power than Deep Blue. Puiu, "Your Smartphone Is Millions of Times More Powerful Than All of NASA's Combined Computing in 1969." See also Nick T., "A Modern Smartphone or a Vintage Supercomputer: Which Is More Powerful?"

64. Nick T., "Modern Smartphone."

65. American Physical Society, "Kasparov Vs. Deep Blue." See also Razmov, "How Computers 'Think' In Chess"; and "Nalimov Tablebases." In chess programming, collections of documented patterns, i.e., chess positions, known as tablebases, are widely used to assess all possible moves and outcomes. For example, the Nalimov tablebase allows perfect calculation of the outcome of any chess endgame with any six pieces or less on the board in any position.

66. Goodfellow, *Deep Learning*, 2.

67. Goodfellow, 2.

68. Singer, *Wired for War.* Examples of previous RMAs in history that commentators offer include the rifle, armored vehicles, aircraft, and nuclear weapons. See also Krishnan, *Killer Robots: Legality and Ethicality of Autonomous Weapons*.

69. Condliffe, "You're Using Neural Networks Every Day Online—Here's How They Work."

70. While the general concepts are consistently mentioned among publications, precise definitions of the related terms do not always enjoy similar consistency, reminiscent of the definitional confusion that has reared its head from time in the cyber community. Perhaps we would be well-advised to heed some of the

years.[71] Although their namesakes suggest a neuroscience and cognitive science basis, it is important to note that AI or learning algorithms may be inspired by biology but also draw inspiration from many fields, especially applied math fundamentals like linear algebra, probability, information theory, and numerical optimization.[72] In other words, they seek to combine the best attributes of human intelligence with the superior computational power of machines.

This concept of intelligence is important to the discussion of AWS because we want any system not under direct control of a human to act, well, intelligently (as we expect humans themselves to act, though often they do not). When we delegate route-finding to our mapping service of choice, we expect it to select the fastest route—not necessarily the shortest one—while accounting for traffic at a given time of day and recent accidents or other delays in determining which route to propose. We expect autonomous cars to be adept at not only following the road and avoiding obstacles but also maintaining real-time situational awareness in cluttered and dynamic environments and exercising appropriate actions for safe driving.[73] And as the debate over AWS highlights, before deploying an AWS that could apply lethal force, we would expect it to be able to reliably identify and discriminate a target as a legitimate military objective and only fire on that target if authorized and if it could do so in a manner consistent with IHL.[74]

In the early years of computer science, the human brain, specifically how it connects millions of neurons to collectively produce intelligent thought, served as an inspiration for neural networks.[75] These neural networks began as mathematical linear models; the models calculated an output from a set of inputs that were weighted to represent their connection to one another and to the ultimate solution.[76] A simple example we can easily imagine is calculating a homework score (the homework "neuron") from the differently weighted inputs of various assignments and then weighting and using that homework score together with other weighted "neurons" (say, class participation, exams,

lessons provided by the growing pains encountered in the cyber domain such as the somewhat confusing terminology often used in discussions on cyber norms. See Osula and Rõigas, *International Cyber Norms: Legal, Policy & Industry Perspectives*.

71. Raghavan et al., *Cognitive Computing*.
72. Goodfellow, *Deep Learning*.
73. Aufrere et al., "Perception for Collision Avoidance and Autonomous Driving." See also Singer, *Wired for War*. Singer discusses some of the challenges facing participants in the DARPA Grand Challenges relating to autonomous vehicles.
74. Anderson, Reisner, and Waxman, "Adapting the Law of Armed Conflict."
75. Goodfellow, *Deep Learning*.
76. Goodfellow.

and so on) to output a final grade for the course.[77] Of course, producing an accurate and reliable output requires that all the weights be set correctly. While the weights in early algorithms were set and adjusted by human operators, over time, systems were developed that could "learn" the right weights on their own by receiving other examples of inputs and expected outputs and then calculating the weights that resulted in the best fitting model.[78]

These early algorithms had many significant limitations, however, in part due to the inherent limitations of the mathematical functions used to model the problems.[79] The decades since have brought about several key concepts that enabled computer scientists to apply this foundational concept of neural networks to more complex challenges, such as image recognition and language processing.[80] Key among these advancements was the concept of depth, as in "deep networks" or "deep learning."[81] Depth refers to the number of levels, or layers, of analysis being done before an algorithm generates the final output.[82] The importance of depth was based on developments in both programming techniques and cognitive science, reflecting the idea that simple computations can be networked to achieve intelligent behavior.[83]

For example, image recognition, say, identifying a face as a face, a relatively simple task for humans, has historically posed a significant challenge for AI systems to perform reliably.[84] However by organizing the problem as stacked layers and combining the results, AI systems are rapidly achieving human or near-human proficiency.[85] Each layer is responsible for analyzing something different. The first might look for a cluster of pixels in an image that resembles an eye.[86] That result is passed up to the next level, which might see if two such clusters were found in relative proximity.[87] The next level might see if that result was associated with a pattern of a nose (which some other "neuron" lower down had already looked for and determined whether a cluster of pixels resembling a nose existed).[88] It might take millions of these "neurons," stacked

77. Though this may be an overly simplistic example, also recall that in the 1940s and 1950s, the state of technology rendered even elementary arithmetic calculations a significant computing challenge; computerized mathematical modeling was thus in fact a notable achievement.
78. Goodfellow, *Deep Learning*.
79. Goldberg, "A Primer on Neural Network Models for Natural Language Processing."
80. Goodfellow, *Deep Learning*.
81. Goodfellow.
82. Goodfellow.
83. Raghavan et al., *Cognitive Computing*.
84. Kaplan, "Machine Learning, Big Understanding."
85. Industrial College of the Armed Forces, *Spring 2011 Industry Study*. In some cases, AI systems are surpassing human ability.
86. Condliffe, "Using Neural Networks Every Day."
87. Condliffe.
88. Condliffe.

up to 15 levels high, to recognize a human face.[89] Accordingly, modern references to "deep" techniques or technologies simply indicate that the algorithm has been designed to derive solutions using multiple layers of computation between the initial inputs and the final output.

Another key advancement was to combine this deeper analysis with improved training of the network to achieve substantial gains in "machine learning."[90] Early training techniques permitted systems to calculate the best weights of the inputs and variables of a particular algorithm through the use of prepared training sets, which are collections of defined input values (i.e., questions and the answer key).[91] These training datasets were often small and narrowly tailored, designed to incur low computational cost and demonstrate that neural networks were able to learn specific kinds of functions.[92] Even today, one of the most important aspects of training any neural network is still feeding the algorithm datasets and then calculating the error between the expected output and the actual output generated by the algorithm.[93] In other words, the algorithm repeatedly engages in trial and error on a massive scale and constantly adjusts the weights to attempt to better align the actual results with the anticipated results.[94] Using our course grade example from earlier, this would be like providing the algorithm the individual assignment and test grades and the expected overall course grades for some number of students and then letting the algorithm calculate the appropriate weights for homework, tests, and so on to reliably determine the right grade when checked against historical data. It can then apply those weights with confidence to calculate future student grades.

What is new about current deeper machine learning is that the drastic increase in computer hardware performance over the last decade in particular and the advent of "big data" have enabled much larger models (i.e., number of "neurons" working together to compute a solution) and much larger datasets to train these models.[95] For example, prior to 2000, datasets may have ranged from the dozens to the thousands of samples, or even less.[96] Toward the end of the first decade of the 2000s, however, and throughout the first half of the 2010s, datasets containing up to tens of millions of examples were produced, including the public Street View House Numbers dataset for learning num-

89. Kelly, "Three Breakthroughs."
90. Goodfellow, *Deep Learning*.
91. Heaton, *Introduction to Neural Networks for Java*, 2nd ed.
92. Goodfellow, *Deep Learning,* 21.
93. Heaton, *Neural Networks for Java,* 119–20.
94. Heaton, 119–20.
95. Goodfellow, *Deep Learning*, 21.
96. Goodfellow, 21.

bers and characters and extensive datasets of translated sentences.[97] These advancements in computational capability and learning have driven consistent improvements in numerous areas such as object and image recognition,[98] speech recognition,[99] language translation,[100] robotics,[101] pedestrian detection and traffic sign classification,[102] to name just a few.

This discussion of algorithms was important to the AWS debate because it reveals why, ultimately, the autonomous trigger pull is not the legal quandary asserted by commentators in this arena. Many of the arguments against AWS appear to overlook or ignore the fact that such systems would be trained, not simply "hard-coded"[103] ahead of time to operate in a rigid and prescriptive manner.[104] Similarly, even proponents of AWS (or at least those not opposed) pay cursory attention to this detail by postulating that concerns with AWS could be mitigated through apportionment of liability to the programmers to ensure they will take greater care[105] or to commanders and other parties who employ the AWS in situations for which it was not designed.[106] A close read of both positions suggests that while they often correctly recognize that AI and machine learning have some role in AWS operation, most claims are nevertheless erroneously based on the concept of deterministic programming: that the AWS must be programmed with a list of if-then propositions and thus might misidentify an input and reflexively perform the incorrect action, or encounter a scenario its programmers did not anticipate and therefore provided no instruction for, leaving the AWS to perform unpredictably or irrationally.[107]

Rather, what we need to understand is that if an AWS ever applies lethal force, it will be because we humans taught the system that in its current situation, given the thousands of conditions and variables being constantly evalu-

97. Goodfellow, *Deep Learning*, 21.
98. Kaplan, "Machine Learning Big Understanding."
99. High, "Era of Cognitive Systems." In 2007 IBM's AI, Watson, appeared on Jeopardy and handily defeated its human opponents.
100. Metz, "Google, Facebook, and Microsoft."
101. Goodell, "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 1." The author describes watching a robot "learn" on its own to walk and to fold towels.
102. Goodfellow, *Deep Learning*.
103. *Oxford Dictionary,* s.v. "Hard-code." Accessed 7 March 2017, en.oxforddictionaries.com. Meaning that rules are not explicitly programmed into the algorithm.
104. Docherty, "Mind the Gap," 8. "Fully autonomous weapons would face great, if not insurmountable, difficulties. . . . It would be nearly impossible to pre-program a machine to handle the infinite number of scenarios it might face."
105. Singer, *Wired for War*.
106. Ohlin, "Combatant's Stance." Ohlin discusses the idea of liability for reckless employment of an AWS as a framework for holding commanders accountable for harm caused by the AWS.
107. Docherty, "Mind the Gap," 9. Reiterates that "due to the infinite number of possible scenarios, robots could not be pre-programmed to handle every specific scenario. In addition, when encountering unforeseen situations, fully autonomous weapons would be prone to carrying out arbitrary killings."

ated[108] and drawing upon the millions, if not tens of millions or more sample inputs and scenarios it has been trained with,[109] lethal force was the appropriate and permissible action to take in accomplishing the mission objectives.[110] More importantly, even if an action results in an undesirable unforeseen outcome, it would still have been a reasonable action for the AWS to take, for the system would have analyzed its vast training data, formulated an understanding of that body of experience, *just like humans do*,[111] and taken action most consistent with that understanding, *which humans often don't*.[112] The difference is that the "decision" of a machine is not tainted by human emotion, such as fear, anger, peer pressure, or even a sense of self-preservation.[113] In other words, an AWS algorithm will always calculate and select the most reasonable action[114] based on significantly more information than a human would possess or be able to process[115] in the heat of battle. If that action in hindsight turns out to have less than desirable effects, there is no reason to believe a human would have fared any better, and quite possibly would have fared even worse if he chose a different, suboptimal course of action. Accordingly, for legal scholars to debate whether employment of AWS is legally permissible due to some inchoate fear or vague notion of unpredictability is misguided. An optimizing algorithm will always select a reasonable action in light of the information available. Therefore, the focus instead should be on how to best collaborate with computer and information scientists, system designers, and developers to articulate and codify the norms we wish to instill in our future systems, much as we expect our leaders to do for young (human) recruits.

---

108. Goodfellow, *Deep Learning*.
109. Goodfellow. Also of note, something that much of the literature seems to acknowledge as a capability but then plays little to no role in the subsequent analysis of legality or propriety is the ongoing training that occurs with intelligent systems. For example, as mature a technology as Google search is, Google still uses about 100 PhD linguists around the world to label and curate training data for its algorithms to ingest and continually refine and improve responses. Similarly, AWS algorithms could be continually trained and refined in response to new environments and scenarios that emerge. *See* Metz, "Google's Hand-Fed AI Now Gives Answers, Not Just Search Results."
110. Although I cannot point to an authoritative source, I would venture that no service member, regardless of training or experience, has ever entered a lethal force scenario with the ability to simultaneously, in real time, evaluate a comparable number of conditions or having the capacity to reference a similar volume of experiential data from which to extrapolate an appropriate response.
111. Feist, *The Psychology of Science and the Origins of the Scientific Mind*.
112. Hammond, Keeney, and Raiffa "The Hidden Traps in Decision Making."
113. Anderson and Waxman, *Why a Ban Won't Work*.
114. The significance of selecting a reasonable action will be explored in greater detail in a later section of this paper.
115. Newman, "Inside Look: The World's Largest Tech Companies Are Making Massive AI Investments." IBM Watson can read 40 million documents in 15 seconds.

Over the River, but Not Yet Out of the Woods

Having defended the proposition that purported legal objections to autonomous systems pulling the trigger are significantly weakened once technically grounded, this section explores why the trigger pull has also been the focus of too much attention in the debate over AWS. A common theme advanced by opponents of AWS is that the threat is imminent and that ceasing or banning AWS research is urgently needed.[116] However this concern that "killer robots" are, literally and figuratively, around the corner overlooks the practical realities of developing and deploying an AWS. These issues are the forest that AWS critics miss for the tree that is the trigger pull.[117] We should instead be turning our attention to a variety of other issues, some that have also been raised in the AWS debate but are substantially satisfied in light of the technical explanations provided in the previous section, and others that are not as prominent but still require resolution and should be addressed in parallel with the associated technological development. A selection of the major issues is presented below.

Autonomy is not a military-specific technology or capability; it has numerous prospects for accomplishing enormous good.[118] Accordingly, investment and advancement in these areas will not be stopped by the current call by some for banning AWS development.[119] This fact highlights that, as alluded to earlier, the term AWS seems a bit of a misnomer. In particular, what is the weapon system that critics decry as being too unpredictable or too ethically or morally repugnant to permit under IHL? Or to borrow an interrogatory from a chapter in the development of law in cyberspace, what is the "thing" that can be reviewed for compliance with international law?[120]

Labeling something as a weapon has far reaching legal, policy, and political implications.[121] It also has practical implications, such as triggering the requirement for legal review as part of the procurement process.[122] Care must be taken then to precisely define or articulate what "thing" constitutes the weapon, lest too narrow or too broad a definition leads to an inaccurate assessment.[123]

116. Gubrud, "Why Should We Ban Autonomous Weapons?" Gubrud argues that "we need to ban [autonomous weapons] as fast and as hard as we possibly can."
117. Calo, "Robotics and the New Cyberlaw." Calo notes "the sorts of problems conscious machines would present are vastly underappreciated." It seems logical that perhaps even without grappling with the thought of machines as conscious, our aperture of the issues presented is still too narrow.
118. Future of Life Institute, "Benefits & Risks of Artificial Intelligence."
119. Ackerman, "We Should Not Ban 'Killer Robots,' and Here's Why."
120. Brown and Metcalf, "Easier Said Than Done: Legal Reviews of Cyber Weapons."
121. Brown and Metcalf, "Easier Said Than Done: Legal Reviews of Cyber Weapons,"128.
122. Brown and Metcalf, 128.
123. Brown and Metcalf, 129.

Take the case of a notional AWS, say, an armed and fully autonomous[124] unmanned aerial vehicle (UAV) that loiters above the battlefield and then lawfully engages predesignated targets or targets of opportunity as they appear. If the UAV was a retrofitted remotely piloted model with autonomy-enabling software installed, would it be overly broad to classify the entire UAV as a weapon system needing review? Would classifying the software collectively as a weapon be overly broad? After all, various software modules control all kinds of benign functions such as takeoff and landing, navigation, environmental sensing, internal system monitoring, and any other number of critical functions one can imagine would be necessary for an aircraft to fly without a human in the pilot seat or remotely controlling its actions. Presumably the actual components that effect lethality, that is, armaments such as missiles, have already undergone their own required legal reviews. Perhaps all that remains is the specific software that performs target selection and engagement. But is that even a weapon? Is it "an instrument of code-borne attack" that could be analyzed to see how its use might comport with IHL,[125] or is the software merely replicating a smart (and presumably law-abiding) human being?[126]

Labeling questions aside, there will assuredly be a legal review involved with at least some component of an AWS.[127] In fact, one of the most developed regimes for testing complex systems in a way that combines law, regulation and technology is the legal review done by DOD lawyers together with technical specialists as part of the weapon system acquisition.[128] The acquisition process thus poses another check on the concerns of AWS detractors. The DOD system acquisition process is the management process by which the DOD provides effective, affordable, and timely systems to the user.[129]

The process begins with an in-depth user need analysis and assessment of existing technology resources and opportunities.[130] Proposed solutions and their capabilities are studied, along with the strategy for developing the neces-

124. Although I use the descriptor "fully autonomous" here to emphasize the hypothetical system capability to operate entirely without human intervention once launched, I concur with the Defense Science Board, "Role of Autonomy in DoD Systems," that "levels" of autonomy (as in not autonomous, semi-autonomous, or fully autonomous) are not helpful to the discussion of AWS and that autonomy should be better viewed as a feature of the capabilities that make up the overall system, as various functions differ in their ability to operate without human guidance.
125. Brown and Metcalf, "Easier Said Than Done," 124.
126. Brown and Metcalf, 124. See also Ohlin, "Combatant's Stance." Ohlin discusses how an AWS might as some point become so sophisticated that it becomes functionally indistinguishable from a human combatant and thus must be treated as one to effectively interact - in opposition or in cooperation - with it.
127. Schmitt, "A Reply to the Critics."
128. Anderson, "Challenges for the U.S. Military in Designing and Deploying Self-Driving Vehicles."
129. DOD, Directive no. 5000.01, *The Defense Acquisition*.
130. *LOG 101 Acquisition Logistics Fundamentals*.

sary technology.[131] Prototypes are developed by competitors in the market, and a final candidate is not selected until after the critical technologies have been developed sufficiently and demonstrated on prototypes in a relevant test environment to satisfactorily reduce the risk of moving forward with procurement.[132] After a winner is selected for formal development, multiple plans are developed, including further technology development, systems engineering, test and evaluation, life cycle support, and others.[133] Once approved, production-representative samples are manufactured in low quantities, demonstrating the manufacturing process and permitting extensive testing in intended environments for assessment against the desired capabilities and technical parameters.[134] Only after passing all these steps and associated requirements and additional subsequent criteria for full production, fielding, testing, and support do the systems get operationally deployed.[135] Furthermore, for an AWS, it is US policy to conduct two legal reviews, once prior to the decision to enter into formal development, and again before an AWS is fielded.[136]

Needless to say, this process is lengthy and detailed. For example, the F-35 Joint Strike Fighter (JSF), arguably one of the most technologically advanced and complex additions to the national arsenal, is not even an autonomous system in the way it has been discussed here, though its software has over 24 million lines of code.[137] In part due to software-based problems causing cost overruns and schedule slippages,[138] the aircraft system has been a work-in-progress for over 20 years and is still undergoing testing and updates.[139] A concern that the DOD is on the cusp of launching an AWS—the technology for which does not yet even exist at the level contemplated or required—into the fray ignores the practical reality of how laborious and test-intensive the acquisition process is intentionally designed to be. Particularly in light of public attention and international scrutiny on the employment of weaponized AI, it would be hard to imagine world-ending AWS as warned about in some of the more sensational claims[140] surviving the DOD acquisition process and associated reviews.

131. *LOG 101*.
132. *LOG 101*.
133. *LOG 101*.
134. *LOG 101*.
135. *LOG 101*.
136. Schmitt, "A Reply to the Critics," 28–29.
137. Charette, "F-35 Program Continues to Struggle with Software."
138. Charette.
139. "F-35 Program Timeline." See also Tucker, "Pentagon Tester: F-35 Program Rushing Tests, Delays Still Likely."
140. Cellan-Jones, "Stephen Hawking Warns Artificial Intelligence Could End Mankind.

Incidentally the procurement process also raises other, though perhaps less spectacular, questions regarding AWS. Federal law, as enacted by the Federal Acquisition Streamlining Act of 1994 (FASA) requires federal agencies to solicit and procure "commercial items" and "nondevelopmental items" "to the maximum extent possible."[141] A "commercial item" is defined by regulation and could be summarized as an item that is available to the public and with reasonable modifications as necessary could be sold or licensed to the government to satisfy a particular need.[142] An inspection of the language in the statute and implementing regulations suggest that where an agency could meet requirements by soliciting and procuring commercial or modified commercially available items (including software products), it must do so.[143] Given the substantial investment of commercial firms in AI and their voracious appetite for limited high-level talent,[144] it is logical that many of the technologies and functions in an AWS will have already been developed for civilian applications and will thus raise issues of commercial item procurement.[145] This is turn may raise potential challenges with intellectual property and staffing or education of sufficiently savvy procurement, legal, and program management personnel.

Beyond domestic concerns, international considerations other than IHL can create operational issues with AWS. For example, South Korea does not allow unmanned aircraft to fly in crowded, unsegregated airspace.[146] As a result, the Northrop Grumman Global Hawk, a bus-sized UAV,[147] must fly hours from its base on Guam to perform reconnaissance missions instead of quickly dispatching from Osan Air Base in the south of the peninsula, as other manned platforms can.[148] Without an agreement or modification to the existing rules, an airborne AWS would likely be subject to the same restrictions. This example represents just one of the legal and practical considerations for AWS employment that will need to be identified and addressed, particularly with allies and partners.

---

141. 10 U.S.C. § 2377 (2011); 41 U.S.C. § 3307 (2010). Applies to both civilian non-defense contracts and to DOD contracts related to national security, defense, or intelligence.
142. 48 C.F.R. § 2.101 (2002).
143. 10 U.S.C. § 2377 (2011); 48 C.F.R. § 10.002(d)(1) (2002); 48 C.F.R. § 11.002(a) (2009); Defense Federal Acquisition Regulations Supplement (DFARS) § 212.212 (2012).
144. Goodell, "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 2." Goodell comments on various events of note in the industry such as Toyota's announcement that it plans to invest $1 billion in a new AI lab and Uber's plunder of the robotics department at Carnegie Mellon University by hiring away 40 researchers and scientists.
145. Winkler, "Palantir Prevails in Lawsuit Over U.S. Army Contracting Practices."
146. Pocock, "U-2 Expert Says Global Hawk Just Can't Compare."
147. Northrup Grumman, "RQ-4 Block 30 Global Hawk."
148. Pocock, "U-2 Expert."

Another post in the fence dividing the AWS debate is the humanitarian aspect. Both Human Rights Watch and the International Committee of the Red Cross (ICRC) have argued that the lethal application of autonomy violates human dignity.[149] This proposition seems premised on the idea that human life should be accorded sufficient respect that the decision to take it away should only be made by a human.[150] Further emphasis is provided with the observation that machines lack qualitative human characteristics such as the ability to perceive subtle behaviors and discern intentions and the ability to identify with other humans.[151] Employment of AWS thus moves humans (and their uniquely human qualities) back one step from the trigger pull.[152]

This step back does not present any new moral issues, however, any more than the morality or lawfulness of ranged weapons, which also precludes the ability to perceive or interact with targets on the receiving end, spurs any serious contention.[153] Rather, where this aspect of the conversation should focus—and where it is resolved—is on the broader question of risk of harm to innocents: civilian bystanders. The taking of an enemy combatant's life, when identified as a lawful military target, must still comply with the rule of proportionality.[154] Proportionality calculations require consideration of both the expected collateral damage and anticipated military advantage; the rule prohibits actions when the collateral damage would be excessive in relation to the concrete and direct military advantage anticipated.[155] In addition to the consideration of the risk of civilian harm incorporated into the proportionality calculation, IHL also requires that attackers take precautions in attack by exercising care to spare civilians and civilian objects.[156]

An AWS can excel at reducing risk of harm to civilians.[157] It can use more capable and precise sensors to more accurately and confidently discriminate

149. Docherty, "Mind the Gap." See also Asaro, "On Banning Autonomous Weapon Systems."
150. Docherty, "Mind the Gap," 9. "As inanimate machines, [AWS] could comprehend neither the value of individual human life nor the significance of its loss."
151. Docherty, 9. Metz, "What the AI Behind AlphaGo Can Teach Us About Being Human." In March of 2016, Google DeepMind's AI named AlphaGo defeated the best Go player in the world 4–1 in a best of five match. This was a milestone in the AI community because over an entire game, the number of possible positions on a Go board is often expressed as exceeding the number of atoms in the universe and so cannot be fully calculated by man or machine. Thus top players play by using intuition, which is typically considered a uniquely human trait.
152. Brown, "Out of the Loop."
153. Brown. See also Schmitt, "A Reply to the Critics," 12; and Ackerman, "We Should Not Ban 'Killer Robots.'" Ackerman discusses how the calculus for employing force has been changing "ever since someone realized that they could throw a rock at someone else instead of walking up and punching them."
154. Schmitt, "A Reply to the Critics," 18.
155. Schmitt, 19.
156. Schmitt, 22. See also Lee, "Double Effect, Double Intention, and Asymmetric Warfare." "Not only should combatants *not try* to harm civilians; they should *try not* to harm them."
157. Which may not be a particularly tall order, at least in the opinion of some. See Arkin, "Warfighting Robots Could Reduce Civilian Casualties, so Calling for a Ban Now Is Premature." Arkin notes that "hu-

potential targets, for example.[158] Moreover, an AWS is not alive and so can be programmed to act more conservatively in self-defense, such as waiting to be actively engaged before returning fire instead of firing first in response to an exhibition of hostile intent.[159] The very lack of human emotion that critics decry as a weakness is in fact a strength when under fire, as these machines will not get tired, stressed, or distracted while waiting or maneuvering to ensure all engagement criteria are met before firing.[160] In fact, some suggest that if the possibility for reducing collateral damage exists, there may even be a moral imperative to use an AWS, similar to the moral imperative advanced by Human Rights Watch to use precision guided munitions in urban settings.[161]

This is not to say that such a system would be easily produced. Roboticists and other experts in the field recognize the challenge of designing a system to effectively model these behaviors and the extensive research that remains to be done.[162] But they take the long view on their research, believing that it could play an important role in reducing noncombatant casualties in future conflicts, the so-called wars after next.[163] In fact, detailed conceptual models have already been developed for technical enforcement of ethical and legal behavior.[164] One proof-of-concept project in particular at the Mobile Robot Laboratory at the Georgia Institute of Technology has been designed, prototyped, and demonstrated as capable of restricting lethal action of an autonomous system in a manner consistent with IHL and rules of engagement (ROE).[165] The demonstration tested only basic scenarios so substantial research and development are still needed, but efforts such as these highlight the promising capability to programmatically govern ethical and legal behavior, squarely addressing a cornerstone of the arguments against AWS.[166]

---

manity has a rather dismal record in ethical behavior in the battlefield."

158. Schmitt, "A Reply to the Critics."

159. Ackerman, "We Should Not Ban 'Killer Robots.' "

160. Ackerman. Cf. Goodell, "Inside the Artificial Intelligence Revolution, Pt. 2." Goodell comments, "I get back in my Hyundai rental after cruising around Mountain View in the [self-driving] Google car, and the first thing I notice is how lousy most human drivers are—pulling out of parking lots without looking, cutting off people during lane changes. I find myself thinking, 'The Google car wouldn't do that.' "

161. Arkin, "Warfighting Robots."

162. Ackerman, "We Should Not Ban 'Killer Robots.' " Ackerman notes that the technology does not exist today nor are the systems ready for fielding in current conflicts.

163. Ackerman. Ackerman interviews roboticist Ronald Arkin from the Mobile Robot Laboratory at the Georgia Institute of Technology, who posits, "I don't believe there is any fundamental scientific limitation to achieving the goal of these machines being able to discriminate better than humans can in the fog of war. . . . [I]f that standard is achieved, it can succeed in reducing noncombatant casualties and thus is a goal worth pursuing in my estimation."

164. Arkin, Ulam, and Duncan, *An Ethical Governor for Constraining Lethal Action in an Autonomous System*.

165. Arkin, Ulam, and Duncan.

166. Docherty, "Mind the Gap."

Perhaps the final major pillar in the case against AWS is the fear that there is no party that might reasonably be held accountable for an egregious violation of IHL, providing an impetus for preemptively banning AWS lest we face such a situation. Undergirding this argument is the fact that an AWS will never "intentionally" violate the law (or "intentionally" perform any action for that matter), as the concept of intent is inapplicable to a machine. As a result, critics assert, the *mens rea* element generally required in criminal violations will never be met, leaving the inadequate option of pursuing liability against a commander, manufacturer, or some other related party.[167] However, recall that, as detailed above, an autonomous system will look across its body of reference data and use thoroughly tested and validated algorithms to always calculate the optimal action to take.

This is key, for in international humanitarian law, the standard is always one of reasonableness.[168] This means acting responsibly when conducting military operations.[169] Attackers must thus consider the information from all sources reasonable available to them at the time, and factors such as force protection, the military value of the target, and the likelihood that subsequent opportunities to conduct an attack will present themselves.[170] It stands to reason that an AWS, with more information, more sensors, and more ability to process that information in the same amount of time is therefore likely to take the most reasonable action. Consequently, an AWS may take an action that in hindsight turns out to be *incorrect*, but it will never be *criminal*.[171]

This truism of course does not relieve the commander of his independent duty to act responsible in employing the forces under his command, including an AWS. This might require steps like ensuring proper maintenance of the AWS prior to use, complying with software patching or update requirements, considering sensor and communication capabilities in the anticipated environment, and performing basic safety and readiness checks as one would normally conduct prior to any other man or machine going "outside the wire" or "wheels up." In addition, blatant breaches such as intentionally programming violative instructions or deliberately bypassing programmatic safeguards and

167. Docherty, "Mind the Gap."
168. Schmitt, "A Reply to the Critics," 21.
169. Schmitt, 16.
170. Schmitt, 16.
171. Schmitt, 21. Schmitt notes how both humans and machines can be mistaken when faced with unexpected or confusing events when making a time sensitive decision in combat but that "[n]either the human nor the machine is held to a standard of perfection" because "the standard is always one of reasonableness." Consider also that humans make decisions of great significance all the time that may be assessed in hindsight as poor, but we excuse or accept as understandable under the circumstances, in fact often criticizing those with too great a proclivity to "Monday morning quarterback" the decisions of others.

ordering the AWS to commit atrocities could be traced back to, and responsibility imposed on, the perpetrators.

In less overt cases of potential wrongdoing, such as negligence or defects or malfunctions of the system, we must remember that the assignment of accountability is not a novel question in the law. Society has already grappled with accountability for commercial products, dangerous animals, slaves and servants, those with diminished capacity, children, employees, soldiers, corporations, and even the occasional unexplainable falling barrel of flour.[172] In each of these cases, the question was not so difficult that it couldn't be answered when new situations emerged, and I see no reason why vigorous discourse and our well-tested system for evolving legal and regulatory controls can't also address future AWS development and use.[173]

Lack of current consensus on legal accountability in various scenarios should not bar concurrent continued work on the significant technical challenges that still must be overcome though. In fact, to place a moratorium or outright ban in the name of first formulating an effective construct for accountability would ironically retard achievement of the very resolution we seek. Without better developed, mature, and robust capabilities, we cannot uncover and appreciate the issues that legal or regulatory mechanisms are best suited to address but remain latent in this nascent stage of autonomy development.

The issue of accountability is closely tied to system auditability. The field of result comprehension has been explored in depth,[174] as it is integral to the establishment of trust in a system.[175] If humans are to pass judgment on the decisions of machines, then a human must have some way of understanding why the machine did what it did—we need an interpretable explanation.[176] This may create difficulties as current testing and validation techniques are generally insufficient for the anticipated complexity of AWS software,[177] and it may not be reasonable to expect a user to comprehend how the hundreds or thousands of inputs contributed to the final action, even if they can review the

172. Asaro, "A Body to Kick, but Still No Soul to Damn: Legal Perspectives on Robotics." "The legal issues raised by robotic technologies touch on a number of significant fundamental issues across far-ranging areas of law. In each of these areas, there can be found existing legal precedents and frameworks which either directly apply to robotics cases, or which might be extended and interpreted in various ways so as to be made applicable." See also Calo, "Robotics and the New Cyberlaw," 129–30; Byrne v. Boadle, 2 H. & C. 722, 159 Eng. Rep. 299 (Exch. 1863).

173. See note 37 above.

174. Ribeiro, Singh, and Guestrin, *"Why Should I Trust You?": Explaining the Predictions of Any Classifier*.

175. Defense Science Board, "Summer Study on Autonomy." See also Endsley, *Autonomous Horizons.*

176. Ribeiro, Singh, and Guestrin, *"Why Should I Trust You?"*

177. Endsley, *Autonomous Horizons.* Verification and validation of advanced systems like AWS will require the development of new methods.

raw data.[178] This remains an open challenge that will need to be addressed so that users will be able to effectively assess the propriety of employment and validate the software's operation of the AWS.[179]

The DOD is sensitive to, and already prophylactically addressing, this human interface issue.[180] In particular, it directs high-level officials to "[d]esign human-machine interfaces for autonomous and semi-autonomous weapon systems to be readily understandable to trained operators, provide traceable feedback on system status, provide traceable feedback on system status, and provide clear procedures for trained operators to activate and deactivate system functions."[181] In addition, these officials must take the necessary steps to "ensure operators and commanders understand the functioning, capabilities, and limitations of a systems autonomy in realistic operational conditions, including as a result of possible adversary action."[182]

Ultimately, ensuring the ability to audit, analyze, and comprehend a system's functioning output is a research, design, and engineering challenge. This challenge does not represent a legal hurdle, and it, too, would be ill-served by a ban or moratorium.

Conclusion

The issues involved with development and employment of AWS are many, but none are unsurmountable. Though some admittedly pose difficult questions, they can be solved with focused attention and informed discussion, more so than with philosophical arguments to remain on the sidelines.[183] However, the long pipeline to a truly effective and ethically governed fully autonomous system is no justification for unnecessary delay or indecision. The threat to national security by not moving swiftly and assertively to maintain at least parity, if not superiority, in autonomy and supporting technologies is very real. The technology is rapidly evolving. China in particular is heavily investing in our domestic technology development,[184] and some sur-

178. Ribeiro, Singh, and Guestrin, "*Why Should I Trust You,*" 188. The authors discuss how the interpretability of explanations must take into account human limitations.
179. Ribiero Singh, and Guestrin. The authors propose a solution that approximates the model that was used by the autonomous system to determine the final action in order to be able to give a simplified explanation to users because the actual calculations would be too complex to understand.
180. DOD, Directive no. 3000.09, encl. 4, para 8.
181. DOD, encl. 4, para 8.
182. DOD, encl. 4, para 8.
183. Brown, "Out of the Loop," 52.
184. Mozur and Perlez, "China Bets on Sensitive U.S. Start-Ups, Worrying the Pentagon." In addition to military hardware like rocket engines and printers that can produce flexible screens for fighter cockpits, China's interest includes sensors for autonomous navy ships, AI, and robotics. One Silicon Valley bank

mise it may soon surpass the United States in AI capabilities.[185] We have the means to invest trillions of dollars in major weapon systems,[186] but AI will enable effects previously achievable only by nation states with entirely low-cost commercially sourced tools.[187] The barrier to entry is simply too low to believe that an international ban or moratorium, much less a domestic instrument implementing the same, will stop development.[188] Autonomy supported by AI and other technologies will only magnify the leverage of potentially adversarial state and nonstate actors against us if we fail to keep pace, and we ignore this possibility at our peril.[189]

To some, this may sound like an AI arms race,[190] and perhaps to some extent it is. But the race is being run whether we participate or not, and the best solution to uncertain or undesirable future events is not to stick one's head in the proverbial sand and hope the worst doesn't come to pass.[191] Rather, we must recognize the urgent need to actively shape and influence the governance of already-proliferating autonomous capabilities. This can be accomplished through focusing and supporting the responsible development of the underlying enabling technology while continuing to define and refine policies, regulation, and legal controls over the eventual products of such technology. The science is taking shape. Conceptual models for ethical governance exist in nascent form. The human capital is available. And as for the lawyers, we should do our part to sharpen the debate and press others to do the same so that when—not if—autonomous weapon systems are fielded by friends and foes alike, we'll be prepared to take a seat at the table in guiding their lawful and ethical use.

---

executive described his interactions with Chinese investors who have "a mandate from Beijing . . . to buy . . . any and all tech."

185. Kania, "China May Soon Surpass America."

186. Francis, "How DOD's $1.5 Trillion F-35 Broke the Air Force."

187. Alexander, "Should AI Be Open?"

188. Scharre, "Ban or No Ban, Hard Questions Remain." "If the technological hurdles are low enough, someone will always cheat."

189. Cf. Freedberg, "Army's Multi-Domain Battle Gains Traction Across Services: The Face of Future War." Army Chief of Staff Gen Mark Milley provides some historical perspective on the need to recognize and incorporate developments that change the way we fight: "We are on the cusp of a fundamental change in the character of warfare. . . . Nations and empires marched off to their destruction [in 1914], blind, *blind* to the changes in war. Let us commit to not march into that abyss, blind to the changes. Let us commit for once, *once* in our history, to not be unprepared for that first battle."

190. Future of Life Institute, "Autonomous Weapons: An Open Letter."

191. Ackerman, "We Should Not Ban 'Killer Robots.'" "Banning the technology is not going to solve the problem if the problem is the willingness of humans to use technology for evil."

Abbreviations

AI	artificial intelligence
AWS	autonomous weapon systems
DARPA	Defense Advanced Research Projects Agency
DFARS	Defense Federal Acquisition Regulation Supplement
DMZ	demilitarized zone
DOD	Department of Defense
FASA	Federal Acquisition Streamlining Act
ICRC	International Committee of the Red Cross
IHL	international humanitarian law
JSF	Joint Strike Fighter
LOAC	law of armed conflict
RMA	revolution in military affairs
ROE	rules of engagement
RPA	remotely piloted aircraft
RUR	Rossum's Universal Robots
UAV	unmanned aerial vehicle

Bibliography

2001: A Space Odyssey. Directed by Stanley Kubrick. Beverly Hills: Metro-Goldwyn-Mayer, 1968.

Ackerman, Evan. "We Should Not Ban 'Killer Robots,' and Here's Why." *IEEE Spectrum*. 29 July 2015. spectrum.ieee.org.

Alexander, Scott. "Should AI Be Open?" Slate Star Codex. 17 December 2015. slatestarcodex.com.

American Physical Society. "This Month in Physics History, February 1996: Kasparov Vs. Deep Blue." *APS News*. February 2002. www.aps.org.

Anderson, Kenneth. "Challenges for the U.S. Military in Designing and Deploying Self-Driving Vehicles." *Lawfare*. 20 October 2016. www.lawfareblog.com.

Anderson, Kenneth, Daniel Reisner, and Matthew Waxman. "Adapting the Law of Armed Conflict to Autonomous Weapon Systems." *International Law Studies* 90 (4 September 2014): 386–411.

Anderson, Kenneth, and Matthew Waxman. *Law and Ethics for Autonomous Weapon Systems: Why a Ban Won't Work and How the Laws of War Can*. Publication. Hoover Institution. 9 April 2013. media.hoover.org.

Andreessen, Marc. "Why Software Is Eating the World." *Wall Street Journal*, 20 August 2011.

Arkin, Ronald C. "Warfighting Robots Could Reduce Civilian Casualties, so Calling for a Ban Now Is Premature." *IEEE Spectrum*, 5 August 2015. spectrum.ieee.org.

Arkin, Ronald C., Patrick Ulam, and Brittany Duncan. "An Ethical Governor for Constraining Lethal Action in an Autonomous System." Technical paper no. GIT-GVU-09-02. Mobile Robot Lab, Georgia Institute of Technology. Atlanta, 2009. www.cc.gatech.edu.

Asaro, Peter. "A Body to Kick, but Still No Soul to Damn: Legal Perspectives on Robotics." In *Robot Ethics: The Ethical and Social Implications of Robotics*, edited by Patrick Lin, Keith Abney, and George Bekey, 169–86. Cambridge, MA: MIT Press, 2011.

———. "On Banning Autonomous Weapon Systems: Human Rights, Automation, and the Dehumanization of Lethal Decision-Making." *International Review of the Red Cross* 94, no. 886 (Summer 2012): 687–709. www.icrc.org.

Attack of the Killer Tomatoes. Directed by John DeBello. San Diego: NAI Entertainment, 1978.

Aufrere, Romuald, Jay Gowdy, Christoph Mertz, Chuck Thorpe, Chieh-Chih Wang, and Teruko Yata. "Perception for Collision Avoidance and Auton-

omous Driving." *Machatronics* 13, no. 10 (December 2003): 1149–161. ri.cmu.edu.
Baig, Edward C. "Robots Will Outnumber Humans in 30 Years, Softbank Says." *USA Today,* 27 February 2017. www.usatoday.com.
Brown, Gary D. "Out of the Loop." *Temple International and Comparative Law Journal* 30, no. 1 (Spring 2016), 43–52.
———. "The Wrong Questions About Cyberspace." *Military Law Review* 217 (Fall 2013): 214–33.
Brown, Gary D., and Andrew O. Metcalf. "Easier Said Than Done: Legal Reviews of Cyber Weapons." *Journal of National Security Law & Policy* 7, no. 1 (February 2014): 115–38. jnslp.com.
Bush, Stephen F., John Hershey, and Kirby Vosburgh. "Brittle System Analysis." Technical paper. Corporate Research and Development, General Electric. 1999. arxiv.org.
Calo, Ryan. "Robotics and the New Cyberlaw." June 2013. robots.law.miami.edu.
Cellan-Jones, Rory. "Stephen Hawking Warns Artificial Intelligence Could End Mankind." BBC News. 2 December 2014. www.bbc.com.
Charette, Robert N. "F-35 Program Continues to Struggle with Software." *IEEE Spectrum*, 19 September 2012. spectrum.ieee.org.
Christian, Brian, and Tom Griffiths. *Algorithms to Live By: The Computer Science of Human Decisions*. New York: Macmillan Publishers, 2016.
Clark, Colin. "Adversaries Outpace US in Cyber War; Acquisition Still Too Slow." Breaking Defense. 19 May 2014. breakingdefense.com.
Condliffe, Jamie. "You're Using Neural Networks Every Day Online—Here's How They Work." Gizmodo. 13 July 2015. Accessed 26 February 2017. gizmodo.com.
"Cyborg / Android / Robot Total Grosses." Box Office Mojo. Accessed 3 March 2017. www.boxofficemojo.com.
DARPA (Defense Advanced Research Projects Agency), "New DARPA Grand Challenge to Focus on Spectrum Collaboration." DARPA. 26 March 2016. www.darpa.mil.
Defense Science Board. "Summer Study on Autonomy." Report. Office of the Secretary of Defense. June 2016. www.acq.osd.mil.
———. "The Role of Autonomy in DoD Systems." Report. Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, July 2012. www.acq.osd.mil.
Department of Defense. Directive no. 3000.09, *Autonomy in Weapon Systems*. 21 November 2012.
———. Directive no. 5000.01, *The Defense Acquisition System*. 20 November 2007.

Dewey, John. *Logic: The Theory of Inquiry*. New York: Henry Holt and Company, 1938.
Dirty. "F-35 Delayed after Fourth Prototype Becomes Self-Aware and Has to Be Destroyed." *Duffelblog*. 4 February 2014. www.duffelblog.com.
Docherty, Bonnie. "Mind the Gap: The Lack of Accountability for Killer Robots." Human Rights Watch. 9 April 2015. www.hrw.org.
Endsley, Mica R. *Autonomous Horizons: System Autonomy in The Air Force—A Path to The Future (AF/ST TR 15-01)*. Report. Washington, DC: United States Air Force Office of the Chief Scientist. June 2015.
"F-35 Program Timeline." Lockheed Martin. Accessed 22 March 2017. www.f35.com.
Feist, Gregory J. *The Psychology of Science and the Origins of the Scientific Mind*. New Haven, CT: Yale University Press, 2008.
Francis, David. "How DOD's $1.5 Trillion F-35 Broke the Air Force." CNBC. 31 July 2014. www.cnbc.com/2014/07/31/how-dods-15-trillion-f-35-broke-the-air-force.html.
Freedberg, Sydney J., Jr. "Army's Multi-Domain Battle Gains Traction Across Services: The Face of Future War." Breaking Defense. 13 March 2017. breakingdefense.com.
———. "Best of 2016: Rise of the Robots." Breaking Defense. 26 December 2016. breakingdefense.com.
———. "F-22, F-35 Outsmart Test Ranges, AWACS." Breaking Defense. 7 November 2016. breakingdefense.com.
———. "Marines Seek to Outnumber Enemies with Robots." Breaking Defense. 25 October 2016. breakingdefense.com.
———. "Should US Unleash War Robots? Frank Kendall Vs. Bob Work, Army." Breaking Defense. 16 August 2016. breakingdefense.com.
Freedberg, Sydney J., Jr., and Colin Clark. "Killer Robots? 'Never,' Defense Secretary Carter Says." Breaking Defense. 15 September 2016. breakingdefense.com.
Future of Life Institute. "Autonomous Weapons: An Open Letter from AI & Robotics Researchers." 28 July 2015. futureoflife.org.
———. "Benefits & Risks of Artificial Intelligence." Accessed 15 March 2017. futureoflife.org.
Goldberg, Yoav. "A Primer on Neural Network Models for Natural Language Processing." *Journal of Artificial Intelligence Research* 57:345–420. November 2016.
Goodell, Jeff. "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 1." *Rolling Stone*, 29 February 2016. www.rollingstone.com.

———. "Inside the Artificial Intelligence Revolution: A Special Report, Pt. 2." *Rolling Stone*, 9 March 2016. www.rollingstone.com.
Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge, MA: MIT Press, 2016. www.deeplearningbook.org.
Goyal, Sanchit. *A Survey on Travelling Salesman Problem*. Department of Computer Science, University of North Dakota. micsymposium.org.
Greenert, Jonathan W. "Kill Chain Approach." *The Official Blog of Chief of Naval Operations*. 23 April 2013. web.archive.org. Web archive version.
Gubrud, Mark. "Why Should We Ban Autonomous Weapons? To Survive." IEEE Spectrum. 1 June 2016. spectrum.ieee.org.
Hammond, John S., Ralph L. Keeney, and Howard Raiffa. "The Hidden Traps in Decision Making." *Harvard Business Review*, 2006. hbr.org.
Harper, John. "Battlefield 2030." *National Defense*, November 2016.
Heaton, Jeff. *Introduction to Neural Networks for Java*. 2nd ed. St. Louis, MO: Heaton Research Inc., 2008.
Hern, Alex. "Give Robots 'Personhood' Status, EU Committee Argues." 12 January 2017. www.theguardian.com.
High, Rob. "The Era of Cognitive Systems: An Inside Look at IBM Watson and How It Works." Technical paper. 2012. www.redbooks.ibm.com.
High, Rob, and Bill Rapp. "Transforming the Way Organizations Think with Cognitive Systems." Technical paper. 2012. www.redbooks.ibm.com.
Holzer, Jenny R., and Franklin L. Moses. "Autonomous Systems in the Intelligence Community: Many Possibilities and Challenges." *Studies in Intelligence* 59, no. 1 (March 2015): 21–29. www.cia.gov.
Industrial College of the Armed Forces. *Spring 2011 Industry Study Final Report: Robotics and Autonomous Systems Industry.* Report. Washington, DC: National Defense University. 2011.
Kania, Elsa. "China May Soon Surpass America on the Artificial Intelligence Battlefield." *The National Interest*. 21 February 2017. nationalinterest.org.
Kaplan, Melanie D. G. "Machine Learning, Big Understanding." *Trajectory*. 3 August 2016. trajectorymagazine.com.
Kelly, Kevin. "The Three Breakthroughs that Have Finally Unleashed AI on the World." *Wired*, 27 October 2014. www.wired.com.
Kinsey, Libby. "A Machine Learning Primer." Medium. 14 April 2016. medium.com.
Krishnan, Armin. *Killer Robots: Legality and Ethicality of Autonomous Weapons*. Farnham, UK: Ashgate Publishing, 2009.
Lee, Steven. "Double Effect, Double Intention, and Asymmetric Warfare." *Journal of Military Ethics* 3, no. 3 (November 2004): 233–51. Accessed 22 February 2017. isme.tamu.edu.

Lewis, Dustin A., Gabriella Blum, and Naz K. Modirzadeh. "War-Algorithm Accountability. Research Briefing. Harvard Law School Program on International Law & Armed Conflict. August 2016. pilac.law.harvard.edu.
Lewis, Harry R., and Christos H. Papadimitriou. *Elements of the Theory of Computation*. 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 1997.
LOG 101 Acquisition Logistics Fundamentals. Defense Acquisition University. icatalog.dau.mil. Course materials on file with author.
Madigan, Tim. "RUR or RU Ain't a Person." *Philosophy Now*, no. 72 (March–April 2009): 48.
Markoff, John, and Matthew Rosenberg. "China's Intelligent Weaponry Gets Smarter." *New York Times*, 3 February 2017. www.nytimes.com.
Martin, David, rep. *60 Minutes*. Season 49, episode 16, "The Coming Swarm." Aired 8 January 2017, on CBS. www.cbsnews.com.
Maslow, Abraham M. *The Psychology of Science*. New York: Joanna Cotler Books, 1966.
"Mathematics and Chess." Chess.com. Accessed 2 March 2017.
McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. Boca Raton, FL: CRC Press, 2004.
Metz, Cade. "Google, Facebook, and Microsoft Are Remaking Themselves Around AI." *Wired*, 21 November 2016. www.wired.com.
———. "Google's Hand-Fed AI Now Gives Answers, Not Just Search Results." *Wired*, 29 November 2016. www.wired.com.
———. "What the AI Behind AlphaGo Can Teach Us About Being Human." *Wired*, 19 May 2016. www.wired.com.
Meyer, Robinson. "The Rise of Progressive 'Fake News' " *The Atlantic*, 3 February 2017. www.theatlantic.com.
Monty Python and the Holy Grail. Directed by Terry Gilliam and Terry Jones. United Kingdom: EMI Films, 1975.
Mozur, Paul, and Jane Perlez. "China Bets on Sensitive U.S. Start-Ups, Worrying the Pentagon." *New York Times*, 22 March 2017. www.nytimes.com.
Murphy, Robin P., and David D. Woods. "Beyond Asimov: The Three Laws of Responsible Robotics." *IEEE Intelligent Systems* 24, no. 4 (July–August 2009).
"Nalimov Tablebases." The Chess Programming Wiki. Accessed 4 March 2017. https://chessprogramming.wikispaces.com/NalimovTablebases (site discontinued).
Newman, Daniel. "Inside Look: The World's Largest Tech Companies Are Making Massive AI Investments." *Forbes*, 17 January 2017. www.forbes.com.
Nick T. "A Modern Smartphone or a Vintage Supercomputer: Which Is More Powerful?" Phone Arena. 14 June 2014. www.phonearena.com.

Night of the Lepus. Directed by William F. Claxton. Beverly Hills, CA: Metro-Goldwyn-Mayer, 1972.
Northrup Grumman. "RQ-4 Block 30 Global Hawk." Accessed 26 March 2017. www.northropgrumman.com.
Ohlin, Jens David. "The Combatant's Stance: Autonomous Weapons on the Battlefield." *International Law Studies* 92, no. 1 (2016). www.hsdl.org.
Orwell, George. *Animal Farm: A Fairy Story*. London: Secker & Warburg, 1945.
Osula, Anna-Maria, and Henry Rõigas, eds. *International Cyber Norms: Legal, Policy and Industry Perspectives*. Publication. NATO Cooperative Cyber Defence Centre of Excellence. Tallinn, Estonia, 2015.
Osula, Anna-Maria, and Henry Rõigas, eds. *International Cyber Norms: Legal, Policy & Industry Perspectives*. Tallinn, Estonia: NATO Cooperative Cyber Defence Centre of Excellence, 2015.
Pocock, Chris. "U-2 Expert Says Global Hawk Just Can't Compare." Breaking Defense. 24 March 2017. breakingdefense.com.
Puiu, Tibi. "Your Smartphone Is Millions of Times More Powerful Than All of NASA's Combined Computing in 1969." ZME Science. 13 October 2015. www.zmescience.com.
Raghavan, Vijay, Venkat Gudivada, Venu Govindaraju, and C. R. Rao. *Cognitive Computing: Theory and Applications*. Amsterdam: Elsevier, 2016.
Razmov, Valentin. "How Computers 'Think' In Chess." Chess.com. 1 March 2010.
Renatus, Flavius Vegetius, and N. P. Milner. *Epitome of Military Science*. Liverpool, UK: Liverpool University Press, 2011.
Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "*Why Should I Trust You?": Explaining the Predictions of Any Classifier*. University of Washington. 9 August 2016. arxiv.org.
Roberts, Adam. *The History of Science Fiction*. London: Palgrave Macmillan, 2006.
"Russia Overtaking US in Cyber-Warfare Capabilities." SC Media UK, 29 October 2015. www.scmagazineuk.com.
Scharre, Paul, and Michael C. Horowitz. "Ban or No Ban, Hard Questions Remain on Autonomous Weapons." IEEE Spectrum, 20 August 2015. spectrum.ieee.org.
Schmitt, Michael N. "Autonomous Weapon Systems and International Humanitarian Law: A Reply to the Critics." *Harvard National Security Journal Feature*. 5 February 2013. harvardnsj.org.
Schuller, Alan L. "At the Crossroads of Control: The Intersection of Artificial Intelligence in Autonomous Weapon Systems with International Human-

itarian Law." *Harvard National Security Journal* 8 (May 2017): 379–425. harvardnsj.org.
Singer, P. W. *Wired for War: The Robotics Revolution and Conflict in the 21st Century*. London: Penguin, 2009.
Slugs. Directed by Juan Piquer Simon. Atlanta: New World Pictures, 1988.
Smith, Chris. *The History of Artificial Intelligence*. University of Washington. History of Computing CSEP 590A. Fall 2006. courses.cs.washington.edu.
Sullins, John P. "When Is a Robot a Moral Agent?" *International Review of Information Ethics* 6 (December 2006): 23–30. www.i-r-i-e.net.
Swarts, Phillip. "RPA Systems Studied to Improve Ground-based Technology." *Air Force Times*, 21 November 2015. www.airforcetimes.com.
The Birds. Directed by Alfred Hitchcock. Universal City, CA: Universal Pictures, 1963.
The Matrix. Directed by The Wachowski Brothers. Burbank, CA: Warner Bros., 1999.
The Terminator. Directed by James Cameron. Los Angeles: Orion Pictures, 1984.
Tucker, Patrick. "Pentagon Tester: F-35 Program Rushing Tests, Delays Still Likely." Defense One. 11 January 2017. www.defenseone.com.
Walsh, Edward J. " 'Robot Ship' Moves Through System Tests." *Proceedings Magazine*, October 2016. www.usni.org.
Williams, Brian, and Emilio Frazzoli. *16.410 Principles of Autonomy and Decision Making*. Massachusetts Institute of Technology. MIT OpenCourseWare. Fall 2010. ocw.mit.edu.
Winkler, Rolfe. "Palantir Prevails in Lawsuit Over U.S. Army Contracting Practices." *Wall Street Journal*, 31 October 2016. www.wsj.com.
Work, Robert O., and Shawn Brimley. *20YY Preparing for War in the Robotic Age*. Publication. Center for a New American Security. 22 January 2014. www.cnas.org.