

The Future of Artificial Intelligence in ISR Operations

COL BRENDAN COOK, RCAF, MSM, CD

Every day, Canada and its allies conduct intelligence, surveillance, and reconnaissance (ISR) operations of one type or another. Despite many successes, operators and analysts have a daily mountain to climb—one which grows with each subsequent mission. That mountain is the result of the continual influx of ISR “big data” that needs to be processed, exploited, and disseminated to end users to ensure the maximum advantage is gained from each mission. Many nations now concede current systems cannot properly analyze and fuse multisensor data. Moreover, these systems cannot provide analysts and operators real-time cues to important information they may be missing. Despite the best efforts to rationalize and realign resources, the mountain of ISR big data grows along with the sense that important intelligence revelations buried in that mountain are being missed.

As with any mountain, there are many paths one can take to the summit. This article aims to chart one path. It will define the ISR community’s big data problem as a way to understand the terrain, explore the potential of artificial intelligence (AI) to address the challenges posed by that terrain, and seek to understand the legal and ethical pitfalls posed by AI. With these factors in mind, this article will present recommendations on how best to develop artificial intelligence, revealing a clear path to the summit of the ISR mountain.

Background

Put simply, AI is a sophisticated decision-making method that enables machines to think and learn on their own.¹ Artificial intelligence differs from autonomy, a broader term referring to “the ability for a machine to perform a task or function on its own.”² Autonomy does not necessarily require AI. In less complex environments, autonomy can be achieved by simple, preprogrammed rules. But more complex, autonomous tasks in open and varying environments do not lend themselves to preprogrammed responses. These tasks require decision-making bordering on cognition—the realm of AI.

Lethal autonomous weapon systems combine autonomy and lethality, may have a human in the loop, on the loop, or human out of the loop, and may or may not possess some form of AI—a feature which often sparks concerns. This article will not address the full breadth of complex problems associated with using these weapon systems. Instead, it will focus on the use of AI in semiautonomous (human-in-the-loop), supervised autonomous (human-on-the-loop), and AI-enabled ISR systems in ISR processes spanning data collection, analysis, and decision-making up to the point of target nomination to a human. In this way, the article will examine what is often considered a less contentious use of AI to determine if some problems and pitfalls remain, even with this limited use.³

Intelligence, surveillance, and reconnaissance is the process by which operators and decision-makers learn about an environment at the tactical, operational, and strategic levels.⁴ Disciples of the revolution in military affairs once preached that the ubiquity of sensing and communications systems would lead to a “powerful synergy” and deliver dominant battlespace knowledge, near-perfect mission assignment, and immediate and complete battlespace assessment.⁵ In an attempt to achieve this vision, militaries worldwide have made significant investments in the ISR enterprise. The Department of Defense (DOD), for example, increased expenditures in ISR systems six-fold from 2001–12.⁶ Similarly, Canada’s latest defence policy leveraged previous commitments and prioritized joint ISR investments to anticipate and better understand potential threats to Canadian interests.⁷ Through these investments, the ISR enterprise now can access data from every domain: air, land, sea, surface, subsurface, space, and cyberspace.⁸ Moreover, the enterprise can draw upon open-source and multilevel classified data.

But the exponential increase in data collection has not led to commensurate improvements in intelligence. As early as 2008, the United States Intelligence Science Board acknowledged that the volume of ISR data exceeded the capacity of the existing analyst community and that much of the data was never reviewed.⁹ In 2014, the RAND Corporation estimated analysts had access to as little as 5 percent of total ISR data.¹⁰ The result for commanders is that fewer intelligence needs are being met.¹¹

To address this deficiency, organizations have improved processes and manning structures, centralizing key functions to maximize manpower, yielding minor improvements in some areas. But the big data problem is only getting worse. Robert Cardillo, director of the National Geospatial-Intelligence Agency since 2014, has noted despite recent improvements, with the current architecture the agency would need 8 million new analysts using current processes to analyze the glut of full-motion video data expected to be collected in the next 20 years.¹² This is but one data source and does not account for the myriad other ISR data sources—

signals, acoustic, radar, and electronic support measures, to name a few—that require analysis to be of any decision-making value.

Challenges

While there is no recognized definition of big data, two recurrent themes emerge: the size and the utility of the dataset. First, big data comprises those datasets “whose size is beyond the ability of typical database software tools to capture, store, manage, and analyze.”¹³ Second, big data consists of information assets whose utility to the organization “demand cost-effective, innovative forms of information processing for enhanced insight and decision making.”¹⁴ These themes are also descriptive. Big data comes from multiple platforms, sensors, systems, and sources that exceed the ability of current database software tools. While highly useful, big data must be given to the right person, at the right time, and in the correct format, to enable decision-making. An analysis of the characteristics of big data sheds further light on why it makes sense to think in terms of ISR big data.

The understanding of the characteristics of big data has evolved. In 2001, Doug Laney proposed the three Ds (data volume, data velocity, and data variety) when analyzing data in e-commerce.¹⁵ A 2014 RAND Corporation study for the US Navy concluded the four Vs can best characterize big data: volume, velocity, variety, and veracity.¹⁶ More recently, other researchers have stressed the importance of adding another V, namely value, to big-data characteristics.¹⁷ These five Vs directly relate to ISR big data.

Intelligence, surveillance, and reconnaissance data is collected in large volumes with a wide variety of formats, sources, and types, and arrives at a high velocity (frequency)—a requirement for delivery to end users.¹⁸ The variety of ISR big data further complicates matters. It may be both open source or classified and, as a result, must be managed across multiple, mutually exclusive security domains.¹⁹ Moreover, ISR big data contains inherent ambiguity, incompleteness, and uncertainty as some data sources are higher quality than others. As such, the veracity of ISR big data must always be challenged and considered when integrating it with other data and information. Lastly, the value of ISR big data is directly related to its role in generating situational awareness and its ability to inform decision-making by being delivered to the right person at the right time and in the correct format.

Opportunities

Having defined and characterized the terrain of the ISR big data mountain, the article will evaluate the promise AI offers to address the five Vs of these data. Since its inception six decades ago, the AI field has alternated between the highs and lows of expectations and actual performance. Setbacks and disappointments have followed periods of great promise.²⁰ The promise has stemmed from the development of AI systems that have progressively challenged humans in gameplay. In 1997, Deep Blue famously beat world chess champion Garry Kasparov, who observed “glimpses of true intelligence and creativity in some of the computer’s moves.”²¹

Advancements since Deep Blue showed promise until recent AI system designs required human intervention to train the systems and necessitated learning from vast amounts of data. Further, these developments demonstrated only a narrow application to gameplay. However, AlphaGo and its successor AlphaGo Zero heralded a new era of AI by demonstrating the ability to play the game of Go, considered the most challenging of human games, at the highest level. AlphaGo was the first AI algorithm to beat human Go champions—the European Champion Fan Hui in October 2015 and Lee Sedol, the winner of 18 international titles, in March 2016.²² In 2017, AlphaGo Zero went one step further, achieving the long-standing goal of learning *tabula rasa* without human intervention. The algorithm learned to play Go through the process, “reinforcement learning, without human data, guidance or domain knowledge,” playing itself in more than 25,000 games.²³ In doing so, the algorithm learned Go from scratch and beat its earlier version 100-0 after only 36 hours of learning.²⁴

While the algorithm’s ability was confined to a narrow task, this experiment demonstrated the potential for AI systems to learn unsupervised. This discovery has opened the way toward artificial general intelligence (AGI), a single system that can learn multiple tasks and employ the knowledge gained in one task to positively transfer over to other tasks—sometimes called meta learning.²⁵ The makers of AlphaGo Zero, DeepMind, announced their subsequent algorithm, Impala, could learn 30 different challenging tasks involving learning, memory, and navigation.²⁶ With AI now on the cusp of AGI, it is poised to provide solutions that will address ISR’s big-data problem.

Artificial intelligence technologies have already been commercialized to address the volume, velocity, and variety of data in multiple fields. The AI employed by John Paul, Amazon, and Netflix have demonstrated the ability to review vast volumes of data regarding customer preferences and available products to provide recommendations for travel needs, online purchases, and entertainment, respec-

tively. Each of these systems analyzes billions of records to suggest products and services based on the previous reactions and choices of users.²⁷

In addition to addressing the volume challenge, economists have turned to AI to address issues of data velocity. Artificial intelligence is being used to create novel data sets from unstructured information, enabling economists to answer questions in real time that previously required months of study. Google has developed systems to analyze search queries to predict changes in unemployment, and Yelp predicts local business patterns, both doing so in real time.²⁸

The ability to process large volumes of data arriving at high velocity is particularly valuable when coupled with AI's ability to analyze many varieties of data such as imagery, speech, language, and electronic signals. Google and Facebook have already deployed face- and image-recognition AI widely in search engines and social media platforms. Project Maven, a DOD initiative, is working with multiple companies to develop image-analysis algorithms to analyze full-motion video data acquired from unmanned aerial vehicles to identify people, vehicles, buildings, and other objects of military value.²⁹ Siri, Alexa, and other personal-assistant AI technologies can already recognize, decode, and translate language.³⁰ The Israeli HARPY missile and US AGM-88 HARM can analyze the radar spectrum, identify enemy radar signatures, and home to targets.³¹

Artificial intelligence architectures have also been proposed and successfully tested to analyze radio signals for a wide variety of applications.³² Each of these specialized capabilities is individually important. A common critique of having specialized AI for each task, however, is that this specialization "inevitably lead[s] to too many network models, increasing the storage complexity."³³ Recent research demonstrated a single AI model constructed from several AI building blocks across multiple domains could be trained concurrently on many data types and tasks.³⁴ Similarly, DeepMind's Impala has demonstrated the capability to conduct many tasks through reinforcement learning. Consequently, AI is already capable of analyzing ISR big data to translate languages, recognize patterns in images and data, find linkages and causation between data, and extract meaning.³⁵ Thus, rather than analysts and operators sifting through raw data, they can now be given the higher-level task of responding to cues, alerts, and conclusions presented to them by an AI-enabled ISR system.³⁶

By fusing and cross-referencing data, these approaches go beyond simply addressing the characteristics of volume, velocity, and variety; they provide mechanisms to address the veracity and value of ISR big data. By overlaying multiple perspectives on each target, the five Vs of ISR big data are satisfied, which improves confidence in the resultant conclusions on target identity, location, motion, and other characteristics. When this process yields conflicting observations, AI

could identify these inconsistencies to operators indicating additional scrutiny is required. Moreover, by providing multiple perspectives on a single target, various low-level features can be extracted and selected from each perspective, and these features can then be compared to identify new, higher-level features in the data.³⁷ Researchers demonstrated this capability by employing a heterogenous, adaptive team of autonomous air and ground robots to monitor a small village; search for, localize, and identify human targets; and simultaneously conduct three-dimensional mapping in an urban setting.³⁸ In these ways, AI systems can be used to ensure the veracity of data while also adding value to it.

Artificial intelligence could also increase the value of ISR big data by alerting analysts and operators to key data and intelligence relating to an area of interest. Siri, Alexa, Google, Amazon, and Netflix AI engines can already monitor user searches and preferences to recommend products and services that anticipate the user's needs.³⁹ Artificial intelligence could monitor the searches and preferences of analysts and recommend data intelligence products to meet their needs. Moreover, as it learns the analyst's requirements, AI could then search through historical data sets to look for patterns of behavior, detect changes, or search for newly assigned priority targets.

For ISR operators, AI algorithms could compare data collected in real time to historical data to ensure sensor operators are alerted to changes from previous observations. Alternatively, as new data from neighboring ISR platforms is collected, it could provide automated cuing regarding observations that may impact the area of operations. These applications would ensure the value of ISR big data is maximized for both analysts and operators, and that less data is lost under the mountain of ISR big data.

A final method AI could use to address the ISR big-data problem is to employ its emerging capacity for creativity, one AlphaGo demonstrated during its second match against Lee Sedol. Midway through this match, AlphaGo made a move that was so unexpected, Sedol paused the game and left the room for 15 minutes to regain his composure. Observers classified the probability that a human would have played that move as 1 in 10,000 and commented that the move displayed "improvisation, creativity, even a kind of grace."⁴⁰ With this level of creativity now possible, AI could be tasked to generate hypotheses about the data it has analyzed. It could then search out data sets to prove or disprove its hypotheses or make recommendations for further ISR data collections. In this way, AI would enable more efficient and focused collections by suggesting collections to prove or disprove its theories, improving the data veracity.

Limitations

Despite the many advantages of employing AI to optimize ISR big data, a question of risk remains. The International Committee of the Red Cross, European Parliament, United Kingdom, the DOD,⁴¹ and others have all considered the implications of employing lethal autonomous weapon systems in warfare. Few have focused on the narrower problem using AI-enabled ISR systems in semiautonomous (human-in-the-loop) or supervised autonomous (human-on-the-loop) modes. But the analysis to date regarding these systems and the work of Nick Bostrom and Paul Scharre regarding risk reduction in autonomous systems, suggest future AI-enabled ISR systems must address the following obstacles: the proper application of the principles of distinction and proportionality; the concerns rising from the “black box” dilemma; the potential for AI systems to mislearn; and the requirement to ensure accountability under the rule of law.⁴²

Under International Humanitarian Law (IHL), the principle of distinction requires attacks only be directed against legitimate military targets. Noncombatants including civilians, children, medical staff, and those combatants considered *d’hors combat*, should be immune from attack, as should civilian objects of no military value.⁴³ To adhere to IHL, an AI system must be able to distinguish between military and civilian targets, a challenge compounded by the fact that no clear criteria exist to make this distinction. It is difficult to instruct or, in the case of AI, to teach a system to avoid targeting civilians and civilian objects when there is no precise specification for “civilians.”⁴⁴

Neither the 1949 Geneva Convention nor the 1977 Protocol 1 define civilian in a negative sense (for example, anyone who is not a combatant) requiring the application of common sense in the determination.⁴⁵ The presence of nonuniformed combatants on the battlefield, particularly in dense urban environments, further complicates matters. Ultimately, an AI system would require a “human understanding of other people’s intentions and their likely behavior” based on subtle cues that may not be easily detectable by sensors or big-data analytics.⁴⁶ In 2013, the Directorate-General for External Policies concluded in its report to the European Parliament that no autonomous system currently exists that can “reliably distinguish between legitimate military targets and civilian persons and objects, [and] take precautions to avoid erroneous targeting.”⁴⁷

More recently, scholarship on the subject concluded that while it may be possible to distinguish cooperative targets that emit known signatures in a controlled environment, accomplishing the same task in an environment with clutter is much more difficult. Moreover, distinguishing an uncooperative target in a cluttered environment is presently beyond the capability of current systems, and “no such

technology is on the horizon.”⁴⁸ Until this challenge can be surmounted, human intervention will be required to ensure the principle of distinction is correctly applied to any targeting decisions.

The principle of proportionality presents another significant challenge for AI systems. This principle requires the expected military advantage to be gained by engaging a target must not be outweighed by the expected civilian collateral damage. While many automated systems can calculate expected civilian collateral damage, there is no objective method to calculate the direct military advantage to be gained.⁴⁹ Absent a method to either program or teach this calculation, there is virtually no way an AI system can comply with this principle on its own. Experts have proposed that human-in-the-loop and on-the-loop autonomous systems do not need to make these judgments on their own to ensure compliance with IHL. By pairing AI systems with humans, the AI system can identify potential military targets and then calculate the potential collateral damage, leaving the human to make the moral judgment.⁵⁰

Beyond the challenges of distinction and proportionality, AI poses a “black box” dilemma. The black box dilemma arises when the complexity in a system increases to the point that a human cannot reasonably understand the process. The human can see the input and output to the system, but the system function is effectively opaque to the user. The principal concern of the black box dilemma is that if a human cannot easily comprehend why and how an AI system is arriving at its conclusions, it is almost impossible for the human to detect when the system fails.

Researchers demonstrated the limitations of the human understanding of AI in a 2013 study of the unexpected outcomes of AI-enabled image identification systems. They studied deep neural networks, a form of AI used in image recognition that had generated counterintuitive conclusions. They found by introducing imperceptible perturbations to images, they could arbitrarily change the AI’s classification of the image. In one experiment, they started with a simple picture of a puppy that was correctly classified by the system. They then made an imperceptible change to the image, only noticeable to the human eye at 10x magnification, and the system then classified the image as an ostrich.⁵¹

Another study investigated this phenomenon from the opposite perspective. The research team trained an AI system to recognize baseballs and then asked it to draw a picture of a baseball. The resulting image was “completely unrecognizable garbage” to a human, but other AI systems agreed with their test system, interpreting the image as a baseball.⁵² Researchers call images that can trick AI systems into misidentifying *adversarial images*. Further work has shown that image recognition software has a widespread vulnerability to adversarial images.⁵³

Consequently, as AI systems develop, humans may not be able to comprehend easily why and how a system arrives at its conclusions.

The difficulties of the black box dilemma can be compounded by the vulnerability of AI systems to mislearn. In March 2016, Microsoft launched Tay on the internet, an AI system designed to exhibit age-appropriate behavior for a teenage girl and to learn through interactions on Twitter.⁵⁴ Microsoft expected Tay to learn millennial slang and start chatting about pop stars. It was instead bombarded with controversial messages from online trolls and within 24 hours was tweeting pro-Nazi messages, denying the Holocaust, and advocating for genocide. Microsoft promptly took Tay offline and issued a formal apology.⁵⁵ This stark example demonstrated the vulnerability of AI systems to mislearn.

The 2016 US election provides a second example where an adversary exploited the use of AI leading to the widespread dissemination of disinformation. As noted in the report to the US Senate, there is compelling evidence that suspected Russian-backed, highly automated, or fake social media accounts were used to sow misinformation and discord in the United States to influence the outcome of the 2016 election.⁵⁶ They achieved this influence by leveraging Facebook, Twitter, Instagram, and YouTube, which each use AI to target users based on interests and behaviors.⁵⁷ In effect, the AI inherent in these social media platforms was exploited to deliver misinformation to American voters on a massive scale. An ISR AI employed to comb through open-source data and classified data in order to deliver useful intelligence to analysts and operators according to their individual preferences could potentially be exploited by an adversary using similar methods.

The vulnerability of AI to mislearn highlights the need to understand AI decision-making with sufficient confidence to ensure accountability under the rule of law. States are obligated under IHL to conduct investigations into the lawfulness of the use of force by their agents.⁵⁸ When incidental civilian death, injury, and/or destruction occurs, or the lawfulness of an attack is in question, an immediate, exhaustive, and impartial investigation must be conducted.⁵⁹ This requirement means information and actions must be traceable in the decision-making process. But if an AI system is effectively a black box—making connections and determinations too complex for any human to comprehend—this becomes problematic, particularly if the AI system cannot be made to explain its reasoning. Therefore, some consideration must be made to ensure some level of transparency exists in an AI-enabled decision-making process to permit detection of failures, prevent mislearning, and for traceability.

Better design, development, testing, and training can minimize the risks of failure in an AI system, but accidents can and will happen with AI-enabled decision-making, just as they do with human decision-making using current

technologies. The accidental shooting down of Iran Air Flight 655 in 1988 by the USS *Vincennes*, and the multiple fratricides by US Patriot missile batteries during the 2003 Iraq War are two examples in which automated systems provided threat indications to operators, who then took what they believed to be an appropriate action.⁶⁰ An AI-enabled system will inevitably result in some failures. Human decision-makers must remain vigilant and closely monitor AI results, with the understanding that this effort may prove difficult on the battlefield.

Experts argue as confidence grows in the use of AI, there is a risk humans will learn to simply trust a system, effectively cease trying to detect failures, and hence become morally disengaged from an AI-enabled decision-making process.⁶¹ There are four known reasons why relying on humans to make decisions based on the assistance of automation can be problematic, each of which played some role in the Iran Air and Patriot missile incidents.

First, reliance on automation leads humans to neglect ambiguity and suppress doubt. Human supervisors then jump to conclusions and cease searching for alternative interpretations to resolve uncertainty.⁶² Second, humans tend to infer and invent causes and intentions by linking fragments of available information through the process of assimilation bias.⁶³ Third, humans are biased to believe and confirm by uncritically accepting suggestions from computers, also known as confirmation or automation bias.⁶⁴ Lastly, a reliance on automation focuses humans on existing evidence and leads them to ignore absent evidence. This phenomenon is often termed “What You See Is All There Is” and “facilitates the feeling of coherence that makes us confident to accept information as true.”⁶⁵ While these factors are all currently at play with existing weapon systems, the black-box nature of AI may magnify these effects, raising the risk humans will cease questioning their “expert AI systems.”

If a human decides on a military action based on the faulty reasoning of an AI system, who is to be held accountable for the decision? There is no easy solution to address this apparent “accountability gap.” Some experts recommend developers pay attention to the human-machine interface design and operator training to ensure that the human-in-the-loop or human-on-the-loop has the capacity and mindset to be responsible for the decisions they make.⁶⁶ Furthermore, AI systems must be designed to allow greater insight into how they arrive at their conclusions and recommendations. Absent these actions, the introduction of AI systems could accelerate existing trends and result in the eventual cessation of effective human supervision.

Recommendations

To summarize, artificial intelligence offers solutions to address the ISR big-data challenge. Well-suited to address the characteristics of ISR big data, the emerging ability of AI to learn without human intervention makes it conducive to manage the myriad of ISR analytical tasks. But the difficulty of providing precise definitions for the principles of distinction and proportionality under IHL will establish an upper limit on what AI can be expected to do. The complexity of AI can render its operation effectively opaque to humans. Adversaries could also leverage the algorithms themselves to disseminate misinformation on a massive scale. The technology is vulnerable to mislearning through the corruption of the data and perverse incentives in algorithms. Moreover, humans are usually predisposed to believe automated systems. All these factors create the risk that humans could become ineffective supervisors of future AI-enabled ISR systems.

To realize the great potential of artificial intelligence and mitigate problems and pitfalls, AI development should be vigorously pursued with four key considerations in mind. First, due to the challenges of defining the principles of distinction and proportionality, there is a limit to the ability of AI technologies to provide highly accurate assessments under realistic combat conditions. Development should be tempered with the expectation that human-machine pairing is both necessary and desirable to ensure compliance with IHL.

Second, the reliance of an AI system on any one source of data to arrive at conclusions may expose these systems to a greater potential to either mislearn or to be manipulated by adversaries. The focus of development should be on building the capacity of AI systems to leverage the volume, velocity, and variety of ISR big data to compare and fuse across multiple data sets. This action will enable the veracity of collected data to be confirmed while simultaneously increasing the value of data and reducing the amount of ISR data left unprocessed and unexploited.

Third, AI algorithms and their associated human-machine interfaces must be designed so that humans can effectively monitor alerts, cues, determinations, and recommendations while also enabling some insight into how AI systems arrive at them. This design would enable humans to detect failures, counter AI's vulnerability to mislearn, and provide transparency during investigations.

Lastly, analysts and operators will require considerable training on AI systems and their employment. This training will need to provide a sufficient understanding of the algorithms to permit the operator to best leverage the potential of AI; methods for the detection of failures and mislearning; an understanding of the potential pitfalls of relying too much on AI and automation in decision-making;

and a recognition of the potential for moral disengagement in AI-enabled decision-making.

With these factors in mind, the potential risks can be reduced and the path AI may offer up the ISR mountain is clearer. The opportunity to choose a better way lies before us. As with all innovations, the implementation of an effective AI-enabled ISR system will take courage, determination, training, and perseverance. Fortunately, these are the same traits that define the modern soldier, sailor, airman, marine, and guardian. The summit is in sight—it is the perfect moment to crest the mountain. ⊕

Col Brendan Cook, RCAF, MSM, CD

Colonel Cook (MSc, Royal Military College of Canada; MS, Air University) is the commander of 14 Wing, Greenwood, Nova Scotia, Canada. He is an air combat systems officer on the CP-140 Aurora (a P-3 derivative) with 30 years of flying experience and an extensive background in underwater acoustic research and air test and evaluation.

Notes

1. Jafar Alzubi, Anand Nayyar, and Akshi Kumar, "Machine Learning from Theory to Algorithms: An Overview," *Journal of Physics Conference Series* 1142, no. 1 (December 5, 2018): 1, <https://iopscience.iop.org/>.
2. Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W. W. Norton, 2018), 27.
3. "Who Decides: Man or Machine?," *Armed Forces Journal*, <http://armedforcesjournal.com/>.
4. B-GA-401-002/FP-001, *Royal Canadian Air Force Doctrine: Intelligence, Surveillance and Reconnaissance* (Trenton, ON: Royal Canadian Air Force Aerospace Warfare Centre, November 2017), <http://www.rcaf-arc.forces.gc.ca/>.
5. Adm Bill Owens, *Lifting the Fog of War* (Baltimore: Johns Hopkins University Press, 2001), 100.
6. Brig Gen Timothy D. Haugh and Lt Col Douglas W. Leonard, "Improving Outcomes: Intelligence, Surveillance, and Reconnaissance Assessment," *Air & Space Power Journal* 31, no. 4 (Winter 2017): 4, <https://www.airuniversity.af.edu/>.
7. Department of National Defence, *Strong Secure Engaged: Canada's Defence Policy* (Ottawa, ON: Department of National Defence, 2017), <https://www.canada.ca/>.
8. Isaac R. Porche et al., eds., *Data Flood: Helping the Navy Address the Rising Tide of Sensor Information* (Santa Monica, CA: RAND Corporation, 2014), <http://www.jstor.org/>.
9. Porche et al., *Data Flood*, 1.
10. Porche et al., *Data Flood*, 14.
11. Haugh and Leonard, "Improving Outcomes," 4.
12. Stew Magnuson, "DoD Making Push to Catch Up on Artificial Intelligence," *National Defense*, June 13, 2017, 22, <https://www.nationaldefensemagazine.org/>.
13. Zhaohao Sun, Kenneth David Strang, and Rongping Li, "Big Data with Ten Big Characteristics," Researchgate, October 2018, 2, <https://www.researchgate.net/>.
14. Porche et al., *Data Flood*, 2.
15. Sun, Strang, and Li, "Big Data."
16. Porche et al., *Data Flood*, 2.'
17. Samuel Fosso Wamba et al., "How 'Big Data' Can Make Big Impact: Findings from a Systematic Review and a Longitudinal Case Study," *International Journal of Production Economics* 165 (July 2015): 234–46, <https://www.sciencedirect.com/>; and Shilian Zheng et al., "Big Data Processing Architecture for Radio Signals Empowered by Deep Learning: Concept, Experiment, Applications and Challenges," *IEEE Access* 6, 2018, 55907–22, <https://ieeexplore.ieee.org/>.
18. Porche et al., *Data Flood*, 2.
19. Porche et al., *Data Flood*, 18.
20. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014), 5.
21. Bostrom, *Superintelligence*, 12.
22. David Silver et al., "Mastering the Game of Go without Human Knowledge," *Nature* 550 (October 2017): 354, <https://www.nature.com/>.
23. Silver et al., "Mastering the Game of Go."
24. Silver et al., "Mastering the Game of Go."
25. Aaron Krumins, "Artificial General Intelligence Is Here, and Impala Is Its Name," *ExtremeTech*, August 21, 2018, <https://www.extremetech.com/>.

26. Krumins, "Artificial General Intelligence."
27. R. L. Adams, "10 Powerful Examples of Artificial Intelligence in Use Today," *Forbes*, January 10, 2017, <https://www.forbes.com/>.
28. Matthew Harding and Jonathan Hersh, "Big Data in Economics," *IZA World of Labor*, (September 2018): 2, <https://wol.iza.org/>.
29. Jon Harper, "Artificial Intelligence to Sort through ISR Data Glut," *National Defense*, January 16, 2018, 34, <https://www.l3harrisgeospatial.com/>.
30. Adams, "Artificial Intelligence."
31. Scharre, *Army of None*, 46–48.
32. Zheng et al., "Big Data Processing Architecture."
33. Zheng et al., "Big Data Processing Architecture."
34. Lukasz Kaiser et al., "One Model to Learn Them All," Cornell University, ArXiv:1706.05137 [Cs, Stat], June 16, 2017, <http://arxiv.org/>.
35. Ethem Alpaydin, *Machine Learning: The New AI*, The MIT Press Essential Knowledge Series (Cambridge, MA: MIT Press, 2016), 55–84.
36. Harper, "Artificial Intelligence," 34.
37. Alpaydin, *Machine Learning*, 74–76.
38. M. Ani Hsieh et al., "Adaptive Teams of Autonomous Aerial and Ground Robots for Situational Awareness," *Journal of Field Robotics* 24, no. 11–12 (November–December 2007), <https://onlinelibrary.wiley.com/>.
39. Adams, "Artificial Intelligence."
40. Joi Ito and Jeff Howe, *Whiplash* (New York: Grand Central, 2016), 240.
41. Noel Sharkey, "Guidelines for the Human Control of Weapons Systems," (International Committee for Robot Arms Control, April 2018), <https://www.icrac.net/>; Nils Melzer, "Human Rights Implications of the Usage of Drones and Unmanned Robots in Warfare," European Parliament Think Tank, May 3, 2013, <http://www.europarl.europa.eu/>; Group of Government Experts of the High Contracting Parties to the Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May be Deemed To Be Excessively Injurious or To Have Indiscriminate Effects, "Human Machine Touchpoints: The United Kingdom's Perspective on Human Control over Weapon Development and Targeting Cycles," August 8, 2018; and Department of Defense (DOD) Directive, *Autonomy in Weapon Systems* (Washington, DC: DOD, May 8, 2017), <http://www.esd.whs.mil/>.
42. Bostrom, *Superintelligence*; and Scharre, *Army of None*.
43. Noel Sharkey, "Saying 'No!' to Lethal Autonomous Targeting," *Journal of Military Ethics* 9, no. 4 (December 2010): 378, <https://www.tandfonline.com/>.
44. Sharkey, "Saying 'No!'," 379.
45. Sharkey, "Saying 'No!'," 379.
46. Sharkey, "Saying 'No!'," 379.
47. Sharkey, "Saying 'No!'," 380.
48. Scharre, *Army of None*, 252–55.
49. Sharkey, "Saying 'No!'," 380.
50. Scharre, *Army of None*, 256–57.
51. Christian Szegedy et al., "Intriguing Properties of Neural Networks," Cornell University, ArXiv: 1312.6199 [Cs], February 19, 2014, <http://arxiv.org/>.
52. Scharre, *Army of None*, 182.

53. Scharre, *Army of None*, 180–83.
54. Luke Dormehl, *Thinking Machines: The Quest for Artificial Intelligence and Where It's Taking Us Next* (New York: TarcherPerigee, 2017), 95–96.
55. Dormehl, *Thinking Machines*.
56. Philip N. Howard et al., “The IRA, Social Media and Political Polarization in the United States, 2012–2018,” Computational Propaganda Resesarch Project, University of Oxford, October 2018, 3, <https://digitalcommons.unl.edu/>.
57. Mike Kaput, “How Facebook Uses Artificial Intelligence and What It Means for Marketers,” Marketing Artificial Intelligence Institute, February 7, 2017, <https://www.marketingai-institute.com/>; and Bernard Marr, “The Amazing Ways Google Uses Deep Learning AI,” *Forbes*, August 8, 2017, <https://www.forbes.com/>.
58. Melzer, “Human Rights Implications,” 40–41.
59. Melzer, “Human Rights Implications,” 40.
60. Scharre, *Army of None*, 137–43, 169–70.
61. Sharkey, “Saying ‘No!’,” 381.
62. Sharkey, “Human Control.”
63. Sharkey, “Human Control.”
64. Sharkey, “Human Control.”
65. Sharkey, “Human Control,” 3.
66. Scharre, *Army of None*, 261.