National Defence Défense nationale

Canadian Forces College

Collège des Forces Canadiennes

MACHINE INTELLIGENCE IN TARGETING: OPPORTUNITIES AND RISKS

Maj Liam Robertson

JCSP 43 DL

Exercise *Solo Flight*

Disclaimer

Opinions expressed remain those of the author and do not represent Department of National Defence or Canadian Forces policy. This paper may not be used without written permission.

PCEMI 43 AD

Exercice *Solo Flight*

Avertissement

Les opinons exprimées n'engagent que leurs auteurs et ne reflètent aucunement des politiques du Ministère de la Défense nationale ou des Forces canadiennes. Ce papier ne peut être reproduit sans autorisation écrite.

Canada

CANADIAN FORCES COLLEGE – COLLÈGE DES FORCES CANADIENNES
JCSP 43 DL – PCEMI 43 AD
2017 – 2018

EXERCISE *SOLO FLIGHT* – EXERCICE *SOLO FLIGHT*

MACHINE INTELLIGENCE IN TARGETING: OPPORTUNITIES AND RISKS

Maj Liam Robertson

Word Count: 3218

Compte de mots: 3218

MACHINE INTELLIGENCE IN TARGETING: OPPORTUNITIES AND RISKS

"Artificial intelligence is the future, not only for Russia, but for all humankind. It comes with colossal opportunities, but also threats that are difficult to predict. Whoever becomes the leader in this sphere will become the ruler of the world.[1]

– Vladimir Putin, 2017

Emerging technologies create conditions for greatly enhanced military capabilities that are potentially disruptive to existing power balances. Chief among the disruptive technologies we face today is Machine Intelligence (MI), also popularly referred to as Artificial Intelligence (AI). The high economic stakes of this technology provide irresistible incentive to compete in what amounts to be an arms race. The Canadian Armed Forces (CAF) and allied forces are investigating the immediate-term challenges and opportunities arising from the implementation of MI to support military operations. The uneven availability of effective MI will challenge our traditional decision-making frameworks and render irrelevant those national capabilities that do not keep pace.

This paper examines the near-term[2] opportunities and challenges related to the implementation of Machine Learning Algorithms (MLA) within CAF Joint Targeting operations, arguing for an investment in specific MI capabilities. It will begin with an examination of considerations for employing MLA in support of military operations and the advantages and disadvantages influencing the transition to MI enhanced warfare. The paper will constrain discussion to the CAF Joint Targeting capabilities emphasizing the near-term human/machine interface opportunities. Finally, recommendations will be proposed for the adaptation of the CAF to MI-enabled warfighting.

---

[1] Vladimir Putin, *Russia Today,* Broadcast national message to students, Moscow, September 1, 2017.
[2] Near-term shall mean deliverable within five years, or out to 2023.

Context

For the purposes of this discussion, the general variety of Artificial Intelligence related terms; Machine Learning, Synthetic Reasoning, Deep Learning Analytics, and Automated Reasoning Capabilities will be referred to as Machine Intelligence (MI). The narrow use of Advanced Neural Networks, Adaptive Algorithms, Machine Reasoning or Artificial Narrow Intelligence to accomplish complex but discrete tasks will referred to as Machine Learning Algorithms (MLA).

The unease created by the potential linkage of Autonomous Weapons Systems (AWS) to MI is not the subject of this paper, however a comment is warranted to preface discussion of MI applications in support of the Joint targeting Cycle (JTC). AWS are not widely fielded at this time, yet the expansion of capabilities in related spheres of endeavour is likely to realize this capability in the very soon. This has motivated eminent figures in the science and technical community to raise the spectre of a dystopian future if unrestrained development and proliferation continue. Machines are already an integral part of warfare, whether they are autonomous or not, and they could undermine global stability. Currently fielded systems permits independent operation within mission parameters, however, a "human in the loop" will not ensure complete safety. The 1988 *USS Vincennes* incident highlights that the human factor is the prime source of erroneous engagement.[3]

[3] LCol David Evans. Vincennes: A Case Study. US Naval Institute, *Proceedings Magazine*, Vol 119/8/1, 086, Aug 1993.

Disruption

Disruptive technical advances permeate our society and rapidly find their way into our future battle-space. The private sector is leading the way in MI research and numerous MLA are available on-line without restriction contributing to a daily trickle of enhancement.

> The holy grail of Silicon Valley entrepreneurs is the disruption of entire industries – because that's where the big money is to be made. Amazon dominates book retailing; Uber decimates taxi services; Pandora replaces radio.[4]

This disruption results in the destruction of livelihoods and shifts in economic and political power, yet the irresistible appeal of large investment returns builds an inexorable force for change.

Human Machine collaboration (sometimes referred to as a "Centaur"), envisions the augmentation of human decision-making rather than replacing it. Human military organizations may be thought of as a complex adaptive system that seeks to dominate other competing complex adaptive systems (our adversaries). The uncertainties surrounding the employment of MLA to support military decision-making make it unlikely that a tightly-coupled system concept will be accepted. Human-machine cooperation in loosely-coupled complex systems is expected to mitigate the inherent risk from common-mode failure of normal accident theory. Future military threats are likely employ larger numbers, and coordinate their attacks more efficiently with more sophisticated maneuver and deception. Therefore military planners and advisors are intent on holding effects delivery authority under human control:

> Saturation attacks from rockets and missiles could overwhelm human operators, a reality that has led over 30 nations to acquire air, rocket, and missile defense

[4] Kaplan, Jerry. *Humans Need Not Apply*, (New Haven: Yale University Press, 2015), 16.

> systems with human-supervised autonomous modes. Future advances in autonomy and swarming are likely to only exacerbate this trend.[5]

Operational time-compression are increasing, however there is no expectation or requirement in the near-term for MI decision-support in the CAF to be directly linked to approving munitions effects.

Near-Term Opportunities

One of the primary means for CAF to contribute to current Coalition operations is by playing role in Joint Targeting campaigns. The Canadian Joint Targeting enterprise has developed recent expertise though limited participation in coalition targeting operations in Afghanistan, Libya, Iraq and Syria. The Chief of Defence Staff (CDS) directed that the CAF establish a comprehensive Joint Targeting capability by Sep 2019.[6] In the near-term, the CAF Joint Targeting enterprise can expect to exploit nascent MLA capabilities in three key areas:

1. Target Development – MLAs support the analysis of large data sets to enable Target Systems Analysis (TSA), Target Audience Analysis (TAA) and the target discovery activities.
2. Dynamic Situational Awareness – MLA agent enhances the authority pathway for Dynamic Targeting and contributes to Common Operating Picture (COP)

[5] Paul Scharre. *Autonomous Weapons and Operational Risk, Ethical Autonomy Project*, (Washington DC, Centre for a New American Security, Feb 2016), 46.

[6] Canadian Armed Forces, CDS Initiating Directive for CAF Joint Targeting, 27 July 2016.

3. Capability/Options Analysis – MLAs rapidly and persistently assess Cyber and Information Domain conditions and propose coordinated courses of action to deliver effects.

Rather than focus solely on MI to enable battlefield automation, the CAF must also ensure that MI enables human cognition, facilitating a competitive JTC capable of delivering full-spectrum effects. Canada must invest in organizations, directed research and designate trial formations to gain understanding of how the integration of MI into military forces will affect future operations. An overview of near-term applications of MLA into the three identified areas of Joint Targeting follows.

Target Development

The detailed staff effort to consider precisely what synchronized effects are desired by a mission Commander requires intensive coordination throughout the targeting cycle.[7] The JTC is an intelligence-enabled activity that is fed vast amounts of data requiring extensive analytical power. At the moment human selected areas are prioritised for examination in detail, leaving much data unexploited. Such target discovery can be more efficiently conducted with MLAs. Canada's space-based Synthetic Aperture Radar (SAR) data is under examination by Convolutional Neural Network to classify images.[8] The resulting analysis has been demonstrated to be highly effective for naval and arctic surveillance exploitation using current sensors.

[7] Canadian and Allied forces stipulate the integrated nature of the Joint Targeting Cycle in various doctrine publications; Canada (CFJP 3-9 Targeting), US (CJCSI 3370.01 Target Development Standards, JP 3-60 Joint Targeting), UK (JSP 900 Targeting Policy).

[8] Ongoing use of MLA to exploit the data feed from the Canadian RADARSAT Constellation Mission and other civilian satellites is encouraging. See Defence Research and Development Canada, *High-Level Information Fusion of SAR and AIS to Enhance Maritime Surveillance,* (Contract Report DRDC-RDDC-2016-C035 November 2015).

The general use of a MLA as an analytical accelerator to classify and match incoming data to indicators and warnings is well understood. The emerging capability to provide conjectures of likely activity based on imagery data is of significant interest to target development organizations. Another example of that capability is the US Algorithmic Warfare Cross-Functional Team (also known as project MAVEN). This project was initiated to accelerate US DoD's integration of large Full Motion Video (FMV) data sets with MLA. The objective was to train the MLA to recognize and cue analysts to potential adversary entities. In late 2017, MAVEN deployed to support operations against the Islamic State.[9]

Dynamic Situational Awareness

As demands for increased precision and tempo continue a cognitive bottleneck obstructs progress. The Intelligent Adaptive Interface (IAI) developments to enhance our dynamic targeting capability utilizing an agent-based algorithm.[10] Enhance decision-making speed with high confidence target sets derived from an extensive list of command options. The potential for this IAI to reduce the workload and human error by efficiently determining the target eligibility for engagement is encouraging. This entails no change in our rigorous target engagement authority process, but will require a cycle of training to engender the trust needed to employ an agent-based Graphic User Interface MLA on operations. A significant key benefit is the requirement to adhere to the established ROE and LOAC in order to continue an engagement. For instance, if Positive Identification (PID) is lost or weapons setting are not consistent with

[9] Kate Conger and Dell Cameron. "Google is Helping the Pentagon Build AI for Drones" (Gizmodo, 6 March 2018).

[10] Dr Ming Hou leads this research at Defence Research and Development Canada (DRDC) Toronto, examining the use of agent-based Intelligent Adaptive Interface (IAI) technology as a decision aid for weapons engagement. See Hou, Ming, A Generic Framework of Intelligent Adaptive Learning Systems: from learning effectiveness to training transfer, *Journal of Theoretical Issues in Ergonomic Science,* Vol 18, 2017.

mission parameters, then weapons release will be denied by the system. The value of an added objective review under combat engagement stress cannot be overstated.[11]

Options Analysis

The leveraging of a Common Operating Picture (COP) that incorporates geographic and temporal representation of all-source intelligence with a natural language generated summary of the data is within grasp. This sort of cognitive persistence is demonstrated with DRDC's WISDOM project. WISDOM is a flexible federation of computer-based tools that support analysts and decision makers in developing their judgement or prediction about situations. It employs MLA to examine a series of related propositions and guide sensor collection. Nested within the DRDC Joint Intelligence Collection and Analysis Capability (JICAC) project, WISDOM has the immediate potential to support options analysis for operational target development.[12] The broad scope of the JICAC project will likely prove unwieldy to the CAF, but the WISDOM architecture is scalable to support current operations. Intelligence augmentation with MLA enables all source intelligence enhancement through document analytics. WISDOM employs multiple automated reasoning tools for multisource exploitation of the massive data repositories analysts are required to sift through to move beyond keyword search and enable contextual searches.

---

[11] Kunduz Hospital was attacked in error with multiple munitions in October 2016. See US Central Command Summary of the Airstrike on the MSF Trauma Center in Kunduz, Afghanistan on October 3, 2015; Investigation and Follow-on Actions 28 April 2016. (accessed 24 May 2018)

[12] DRDC Project 05da: JICAC is a series of experiments and trials that explore the integration of a number of national intelligence collection and analysis activities. For insight into one of the see Shadi Ghajar-Khosravi and Peter Kwantes, "Sharik 2.0: The Design and Development of a Web-Based Tool to Support Collaborative Sensemaking", *Scientific Report*, (DRDC-RDDC-2017-R106, August 2017).

Proposition queries to explore database and assist analysts in understanding a target system or audience. A synthesis of human developed engagement options or courses of action (COAs) contrasted with computer-generated COAs becomes possible. An MLA observes information differently than its human counterparts and is capable of identifying patterns in the data that human analysts may overlook due to the large volume and complexity of data. The potential for the further exploration of imperfect information is illustrated with the Libratus poker system which seemingly addresses a challenge in game-theoretic reasoning containing hidden information within a large state space.

> The techniques that we developed are largely domain independent and can thus be applied to other strategic imperfect-information interactions, including non-recreational applications. Owing to the ubiquity of hidden information in real-world strategic interactions, we believe the paradigm introduced in Libratus will be important for the future growth and widespread application of AI.[13]

Such an MLA can also be employed to enhance the delivery of effects in the Information Environment (IE). State and non-state actors can leverage greater sophistication and scale in executing their Information Operations.[14] The rapid creation of precise messages and media formats by MLA enables the delivery of rapidly adaptive synchronized non-munitions effects across all contested domain.

Preliminary studies also indicate that MLA-enabled offensive cyber operations will significantly enhance cyber effects[15]. The addition of MLA to a cyber operation will increase the tempo and variety of cyber effects while reducing the cost. Proliferation of specialized MLA will

---

[13] Noam Brown, and Tuomas Sandholm, "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals", (Science Magazine, American Association for the Advancement of Science, 17 Dec 2017), 7.

[14] Brundage, Miles, *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*. (Future of Humanity Institute, University of Oxford, 2018), 45.

[15] Ibid. 34.

enable a broad swath of actors to engage in this activity. The key limitation will become access to the suitable algorithms and the requirement for skilled specialists will diminish.

Risks

MI is quintessentially dual use in nature. This ensures that discoveries that may yield commercial benefit will proliferate rapidly. Once proven and trusted, an MLA is materially more efficient than human centric processes and vastly scalable. The classic business objective of "faster, better, cheaper" is within grasp. Any attempt to constrain or ban such technologies will be fraught with difficulties. A motivated actor with the resources to purchase or steal a MI system may obtain tremendous offensive cyber-capability, potentially becoming an Advanced Persistent Threat (APT) even if that actor is relatively ignorant of the technology. The marginal cost of replicating software approaches zero, providing no constraint. MLAs like this could then be readily adapted to drive APT cyber-attack tactics, where the MLA is competing against human or non-adaptive defensive MLA.

MLA employing deep neural networks have proven to be an extremely powerful tool for object recognition, often can performing as well or better than humans in standard testing. Despite that success, some unexpected vulnerabilities persist. A class of visual objects known as "adversarial images" are capable of deceiving algorithms into identifying false images with high levels of confidence.

> This vulnerability of deep neural nets to adversarial images is a major problem. In the near term, it casts doubt on the wisdom of using the current class of visual object recognition AIs for military applications – or for that matter any high-risk applications in adversarial environments.[16]

[16] Paul Scharre. *Army of None: Autonomous Weapons and the Future of War*, (New York, W.W. Norton & Company, 2018), 182.

This subtle manipulation of imagery data (data poisoning) or adversarial inputs reveal an emergent phenomena stemming from fundamental properties of the internal structure of neural networks. An educational loop to train the algorithm against such poisoning is ineffective since the space of all possible images is effectively infinite and the system can fail in ways humans would not. Even without knowing how a specific neural network is structured, an adversary could generate deceptive images in various media, creating false target indicators and concealing actual entities. This latent ability for exploitation of design flaws lays bare the essential requirement for human-machine cooperation in target development activities.

Due to the ability to exceed human capabilities, it has impact on human community, much like GPS has done to human navigation skills. Extended use of MLA in any area of endeavor may increase the psychological distance operators have with people they are planning to influence. The popular devotion to social networking, on-line gaming and internet surfing inculcates a potential for a reduced capacity to empathize with others.[17] This factor has a potential to bias human analytical processes and will require a conscious effort of leaders to regulate.

The rapid adoption of MLA and other disruptive technologies require concurrent operational adaptation. Such transformation requires our forces to re-tool and conduct comprehensive capability analysis. Upgraded training and exercises will be needed for forces to work seamlessly with MLA and other systems which consumes significant investment and management resources. With the MI arms race underway, fear of having one's national project

[17] Greenfield, Susan. *Mind Change: How Digital Technologies Are Leaving Their Mark On Our Brains,* (New York: Random House, 2014), 118.

overtaken encourages the technological competition and a race to the bottom.[18] These risks reinforce the fundamental theme that we cannot place trust in MLA or MI systems to be infallible, such systems will require continuous in-depth validation reviews and will still require supervision.

CAF Adaptation

The natural fossilization of institutional thought limits how quickly the CAF can perceive and explore emerging concepts. We possess well-honed industrial age institutions that are struggling to remain relevant in the information age. The CAF does not have the scale to create leading edge capabilities, but we are small and agile enough to rapidly operationalize functional concepts in order to seize relevance in key areas. In this manner the CAF can be "Future-proofed" to the MLA capability gap by establishing an internal centre of excellence for MI. Like many other organizations a culture of continuous improvement is required just to maintain our current relative position, not to mention establishing competitive advantage.

The cornerstone to such future-proofing is the establishment of a multidisciplinary cadre of capable CAF and DRDC members possessing MI/MLA skills and capabilities. Such a community of interest could then be provided venues to conduct formal gap analysis and brainstorm the options for CAF experimentation in this area. This falls under the broad scope of responsibility of the Chief of Force Development (CFD), where MI is understandably difficult prioritize. The potential for collaboration with the Canadian commercial sector is growing and this may well be where the momentum for institutional involvement will originate.

---

[18] Nick Bostrom. *Superintelligence: Paths, Dangers, Strategies*. (Oxford: Oxford University Press, 2014), 247.

Once the MLAs are tailored and implemented for the aforementioned three targeting areas of opportunity attains full operating capability, it requires validation. The conduct of challenging and realistic wargames is essential to understand and socialize the use of these improved capabilities. Events such as the Schriever Wargame 2017 rigorously challenge capabilities that are anticipated to be in operation ten years in the future.[19] Integration of our targeting enterprise with our key allies is essential for Canada top participate in future joint coalition operations.

As this technology advances, more MLA will be available for a variety of activities, particularly for the delivery of malicious software. The development of tactics and preparation of attacks still require human expertise for the foreseeable future. The aforementioned impact of MLA accelerates the complexity and velocity of change facing military planners. The central role of the human is perceived to be under threat and there will be institutional resistance to change. However, humans remain essential:

> These and similar errors are often classified as "human errors': it wasn't the system that was at fault; it as the programmer, engineer, or user who did something wrong. But it might be fairer to call them "human to computer translation errors": a human does something that would make sense if they were interacting with another human, but it doesn't make sense to a computer.[20]

Humans are an adaptive species and the integration of MLA within all aspects of civilian society will have transformational effects on global culture. The CAF will need to invest in systems and infrastructure that are capable of running and sustaining the increased computational power that comes with training and deploying MLA. This requires increased ties with

[19] The Schriever Wargame is an annual event that examines a near-future peer adversary conflict employing capabilities that are expected to be mature within 10 years. See US Air Force Space Command, *Schriever Wargame Concludes*, 20 Oct 2017, 2018http://www.afspc.af.mil/News/Article-Display/Article/1349906/schriever-wargame-concludes/ (accessed 12 May 2018)

[20] Stuart Armstrong, *Smarter Than Us: The Rise of Machine Intelligence*. (Machine Intelligence Research Institute, Berkeley, 2014), 18.

commercial and academic institutions to keep CAF deployed hardware and software on the leading edge of the global paradigm.

Looking forward

The private sector is leading the way in MI research and Canadian world-class expertise is actively recruited by US corporations. Albeit certain elements within the technology sector appear reluctant to partner with defence or security agencies. Some entities going as far as publicly stating policies that they will not contract with defence or security agencies. This challenge has been effectively taken up by Canadian national leadership and CAF/DRDC researchers are regularly invited to MI technical seminars, conferences and trade shows.

CAF does not have the scale to create leading edge capabilities, but we are small and agile enough to rapidly operationalize functional concepts in order to seize relevance in key areas. The CAF may be able to Future-proof the MLA capability gap by establishing a narrow internal centre of excellence for MI. In the long term, the establishment of a strategic Nash equilibrium in the MI realm may be the global transition solution, a condition where no powerful MI entity has anything to gain by changing their strategy while the other entities keep theirs unchanged.

Ideally Canada should implement a national programme to guide and coordinate effort along the lines of the Israeli Technion Institute cross-sector interdisciplinary research approach that contributes to worldwide technical development.[21] Perhaps we need to secure access rights to expertise, technology and data through national legislation or some other means. It is also

[21] Technion, Israel Institute of Technology is a comprehensive academic center for advanced science education and applied interdisciplinary technical research. See www.technion.ac.il.

worth considering the establishment of a register of security cleared Canadian nationals possessing MI/MLA capabilities to be called upon in times of national emergency.

Somewhere, modern digitally-native youth infused with a culture of innovation and risk-taking will seize and retain the technological initiative. MI systems are driving global economic and military innovations and Western nations must stay in the lead. The combination of people, training, doctrine, experimentation and validation trials is our key advantage and difficult to emulate. Canada has the pieces, but we are not coherently organized to move forward at the moment. Our emerging rivals are well aware of the advantages provided by MI and must be anticipated they will employ all available options to counter our strength in the traditional domains.

Conclusion

The foregoing discussion has asserted that near-term conflicts will involve the use of MLA to support the planning and delivery of cross-domain effects. The ability to produce synchronized rapidly adapting multi-domain effects will become ever more reliant on MI, even if it is tightly paired with human guidance. The *sine qua non* of victory will be the possession of capable Machine Intelligence. In the near term, investment in such capabilities by our rivals will only increase.

Future victory will still be achieved in the Human/Informational domain and it will belong to the humans who retain the responsibility for the effects of the operations that they plan and execute. This paper has examined the immediate-term opportunities and challenges arising from the implementation of MLA on the CAF Joint Targeting enterprise arguing for an investment in specific MI capabilities. It also identified some of the key factors for employing

MLA in support of military operations and the advantages and disadvantages influencing the transition to MI enhanced capabilities. As the pace of warfare accelerates and trust in automated systems matures, an adaptation of human military culture must occur for the CAF to remain competitive.

In summary, the following CAF investment in emerging MI capabilities and select MLA is required in order to maintain our military value:

- Integrate the nascent MLA capabilities into the current CAF Joint Targeting enterprise (Target Development, Dynamic SA, and Options Analysis) and collaborate with allies in maintaining MI/MLA dominance;
- Future-proof the CAF to the MLA capability gap by establishing an internal centre of excellence for MI;
- Partner with select Canadian commercial and academic entities;

The high economic stakes provide irresistible incentive to compete in what amounts to be a technical arms race. Our rivals are not shrinking from the challenge. President Putin's national education message was a raw challenge to the youth of Russia to become a leader in AI; we must not allow our innate Western hubris to ignore this challenge.

BIBLIOGRAPHY

Armstrong, Stuart, Smarter Than Us: The Rise of Machine Intelligence. Machine Intelligence Research Institute, Berkeley, 2014.

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.

Brown, Noam, and Sandholm, Tuomas, "Superhuman AI for heads-up no-limit poker: Libratus beats top professionals", *Science Magazine*, American Association for the Advancement of Science, 17 Dec 2017.

Brundage, Miles, *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*, Future of Humanity Institute, University of Oxford, 2018.

Canada, Canadian Forces Joint Publication, 3-9, Targeting, 2014.

Canadian Armed Forces, Chief of Defence Staff Initiating Directive for CAF Joint Targeting, 27 July 2016.

Conger, Kate and Cameron, Dell. "Google is Helping the Pentagon Build AI for Drones" *Gizmodo*, 6 March 2018. (accessed 05 April 2018)

Defence Research & Development Canada, Joint Intelligence Collection and Analysis Capability (JICAC) Project, April 2018.

Evans, LCol David. "Vincennes: A Case Study". *Proceedings Magazine*, Vol 119/8/1, 086, US Naval Institute, Aug 1993.

Greenfield, Susan. *Mind Change: How Digital Technologies Are Leaving Their Mark On Our Brains*, New York: Random House, 2014.

Hou, Ming. "A Generic Framework of Intelligent Adaptive Learning Systems: from learning effectiveness to training transfer", *Journal of Theoretical Issues in Ergonomic Science*, Vol 18, 2017.

Kaplan, Jerry. *Humans Need Not Apply*, New Haven: Yale University Press, 2015.

Putin, Vladimir. *Russia Today*, Broadcast national message to students, Moscow, September 1, 2017. Translated by Russia Insight youtube channel 4 Sep 2017 (accessed 24 May 2018)

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

Scharre, Paul. *Autonomous Weapons and Operational Risk*, Ethical Autonomy Project, Washington DC, Centre for a New American Security, Feb 2016. (accessed 12 April, 2018)

USDI Algorithmic Warfare Cross Functional Team, Project MAVEN update brief, April 2018.

US Air Force Space Command, *Schriever Wargame Concludes*, 20 Oct 2017, 2018http://www.afspc.af.mil/News/Article-Display/Article/1349906/schriever-wargame-concludes/ (accessed 12 May 2018)

US Central Command, *Summary of the Airstrike on the MSF Trauma Center in Kunduz, Afghanistan on October 3, 2015; Investigation and Follow-on Actions,* 28 April 2016. (accessed 24 May 2018)