Simplifying Emerging Technologies

Risks and How to Mitigate Them

Edited by Kristen Csenkey,
Balsillie School of International Affairs,
Waterloo

May 2020

Supported by the Department of National Defence Mobilizing Insights in Defence and Security
Targeted Engagement Grant

Contents

Bridging the Gaps: An Introduction to the Risks and Opportunities of Emerging Technologies

Kristen Csenkey, Balsillie School of International Affairs, Waterloo

Introduction

Quantum computing, artificial intelligence (AI), encryption, and information and communication technologies (ICTs), provide Canada with possibilities – to process data faster, to develop a competitive edge against adversaries, and to make Canadians safer – but there are also risks. According to Strong, Secure, Engaged (SSE), adversarial actors can use emerging technologies to engage in threatening behavior faster, easier, and more conveniently than ever before. These threats can take shape through attacks on critical infrastructure, the democratic process, and Canadian data. Canada must adapt to these threats, but how can we adapt unless we understand the challenges involved?

To answer this question, this publication addresses these challenges by 1) identifying evolving emerging technologies, 2) explaining their utility, 3) discussing their practical, legal, and ethical implications in a military context, and 4) providing recommendations based on these implications.

SSE calls for the Department of National Defence and the Canadian Armed Forces (DND/ CAF) to adapt to the rapidly evolving defence and security landscape. Dual-use and emerging technologies are a part of this dynamic landscape, and the authors within this publication show how Canada can navigate this terrain. The authors conceptualize the threats and opportunities of emerging technologies around four themes:

- data and network security,
- governance of emerging technologies,
- maintaining national sovereignty of innovation; and
- AI and decision-making.

These themes show that the development and use of emerging technologies are increasingly complex and are found within multiple domains of engagement. Dr. Nina Bindel's piece, "Ensuring

Data and Network Security in the Era of Quantum Computers" focuses on addressing the first theme with tangible examples and explanations. She shows the risks and benefits associated with quantum computing and argues for quantum secure alternatives. In "Governance of Emerging Technologies: Canadian Cyberspace Governance — or Lack Thereof?", Josh Gold calls for the increased transparency and action on an international stage. Dr. Travis Morrison argues for the protection of Canadian data in his brief, "Homomorphic Encryption and Secure Outsourced Computation". In the final piece, "Human-Machine Teams in Near-Future Military Environments" by Dr. Sarah Shoker, discusses the challenges of AI decision-making in conflict situations.

Using emerging technologies to advance Canadian security objectives and strategic interests is a two-part mission. The first part is ensuring clarity. Those with technical knowledge can help articulate the details of these technologies. With this information in hand, we can then move to discuss the practical implications, including the associated risks and opportunities. The fog of uncertainty and misunderstanding about the capabilities of quantum computing, AI, encryption, ICTs, and AI needs to clear. This project helps us achieve this goal.

Acknowledgements

This endeavour was made possible through the Mobilizing Insights in Defence and Security (MINDS) Targeted Engagement Grant. Through this grant, the Department of National Defence invested in academic research by young professionals and post-doctoral researchers. It brought together diverse perspectives from quantum computing, cryptography, privacy, and international relations to identify and attempt to solve the critical issues associated with emerging technologies and their military applications.

This publication would not be possible without the dedication and support of the following people. Aaron Shull (Centre for International Governance Innovation) and Dr. Douglas Stebila (University of Waterloo) took the time to review and provided valuable feedback on the briefs. Josh Gold and Drs. Nina Bindel, Travis Morrison, and Sarah Shoker provided their expertise on the topic of emerging technologies. The view expressed in this document based on the authors' own opinions and research. Nathan Grift created the infographics based on each of the briefs. Chris Earle, Kersty Kearney, Nawroos Shibli, and Sarah Wyatt worked as part of the research team to pull all the threads of this project together. The biographies of the contributing authors and research team members are located at the end of this document.

The editor and authors are grateful for the support from the Department of National Defence and the Canadian Armed Forces, the Balsillie School of International Affairs, the research team, and to everyone who participated in the making of this publication.

Ensuring Data and Network Security in the Era of Quantum Computers

Dr. Nina Bindel, University of Waterloo and Institute of Quantum Computing

Introduction

One of the biggest challenges regarding emerging technologies is that the risks are acknowledged and, often, actions are taken too late. When implementing risk prevention, the difficulty is then to balance acting hastly and acting cautiously regarding decisions, recommendations, and regulations. This is particularly true for decisions about how to ensure security against quantum computers—decisions that must be made now.

The Quantum Computer Threat

Quantum computing and its applications, such as large-scale quantum computers, promise to enable computations that are otherwise too inefficient for current computers. Recent advances, e.g. the proof of quantum supremacy,[1] are major steps towards building large-scale quantum computers. However, these advances also pose a threat: using large-scale quantum computers will make it possible to break essentially all public-key cryptography in-use today.

> *“As long as the process of a foreign IT-security standard has been transparent, it is reasonable to adopt it, given the urgency to act.”*

The reason is that the security of standardized (asymmetric) cryptographic algorithms that are in use today, such as FIPS 186-4 for digital signature schemes[2] or RFC 8017 for digital signature and public-key encryption schemes (PKEs),[3] are based on the difficulty of the integer factorization or the discrete logarithm problem. While we are not aware of any algorithm that would break these

[1] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (2019), pp. 505–510.

[2] National Institute of Standards and Technology. *Digital Signature Standard (DSS)*. 2013. URL: https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf.

[3] Kathleen Moriarty et al. “PKCS# 1: RSA cryptography specifications version 2.2”. In: *RFC 8017* 10 (2016).

mathematical problems in polynomial time on current (also-called classical) computers, Shor's quantum algorithm[4] solves these problems in polynomial time using quantum computing. Hence, as soon as quantum computers exist that are large enough to implement and run Shor's algorithm for key sizes in-use, e.g. 2048-bit RSA keys, our security guarantees do not hold anymore. Given the importance of IT-security ensured by cryptographic algorithms, this would have a serious impact on the safety and economic well-being of ordinary people as well as companies and governments.

Quantum Secure Alternatives

To prepare for this security threat, cryptographers have been working on alternative algorithms which are not known to be vulnerable to quantum attacks—the so-called post-quantum or quantum-secure cryptographic algorithms. In order to advance this effort, in 2017 the US-American National Institute of Standards and Technology (NIST) launched a new standardization process with the goal of selecting the next generation of quantum-secure public-key cryptographic algorithms.[5] The initial phase of this process ended with the selection of 17 key encapsulation mechanisms (KEMs) or PKEs and nine digital signature schemes for the second phase in March 2019. NIST is planning to announce the candidates for the third round in June 2020. According to the current schedule, the final standards will be available in 2022/2024.

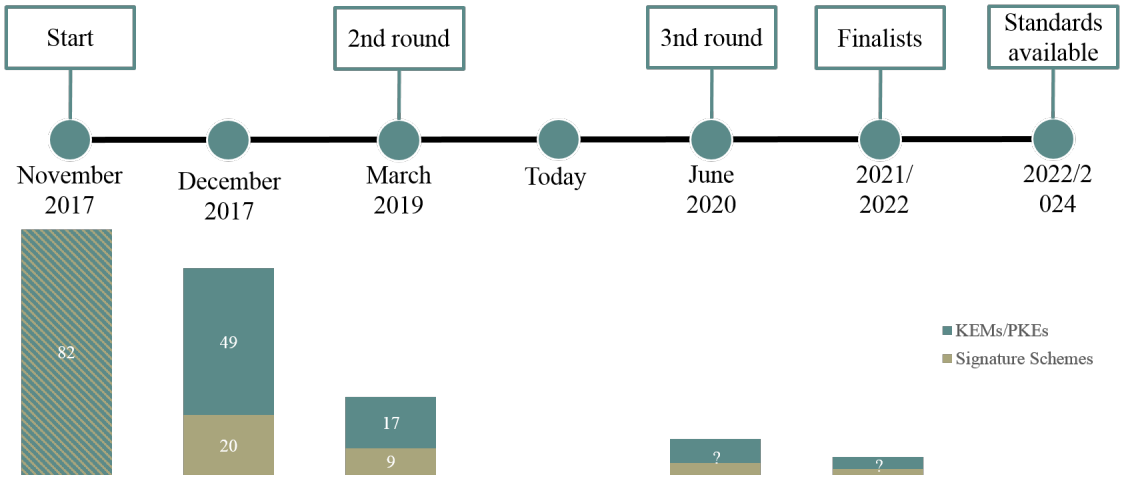

Figure 1: Timeline and number of candidates of NIST's standardization effort

Canada's Sovereignty vs. Its Partnership with NIST

Indisputably the outcome of NIST's post-quantum standardization effort will impact the decisions of other standardization agencies worldwide. In particular, close partners of the U.S, such as Canada,

[4]Peter W Shor. "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer". In: *SIAM review* 41.2 (1997), pp. 303–332.

[5]National Institute of Standards and Technology (NIST). *Post-Quantum Cryptography Standardization*. 2017. URL: https://csrc.nist.gov/Projects/Post-Quantum-Cryptography.

are likely to favor the same algorithms. This is particularly reasonable in this case as researchers recommend a transition to post-quantum cryptography sooner rather than later.[6] However, since earlier standardizations gave reason to be cautious[7], no standardization agency should blindly follow the decision of another country.

NIST's post-quantum standardization effort, however, seems to be trustworthy for the following reasons. It is community-based: from the very beginning NIST asked for public comments on the call for proposals as well as the algorithms using a comment function on their website. This increases the chance to detect weaknesses. Furthermore, third-party projects to evaluate the schemes seem to be taken into consideration. Prominent examples are the following benchmarking and testing platforms: SUPERCOP, SafeCrypto/pqclounge, pqm4, or the Canadian framework liboqs.[8] Moreover, the NIST standardization effort is very international with 69 researchers from different countries, including researchers affiliated to Canadian universities and/or corporations. This diverse expertise is beneficial to foresee different kinds of risks. Furthermore, NIST allows a rather large degree of flexibility in that changes to the submissions are encouraged in order to improve the algorithms.

Another important property to ensure trust in a standardization process is transparency. We will have to wait and see how transparent NIST's final decision will be. In particular, before following NIST's standardization, one should ask the following questions:

- What is the reason why a particular candidate was preferred over another one with similar properties? This could give insight into non-scientific selection biases.
- Is every change between the submissions and the final standard explained and evaluated by NIST, the researchers, and/or the community? This concerns the choice of new parameters as well as implementation changes in order to avoid backdoors or other similar risks.

As long as the process of a foreign IT-security standard has been transparent, it is reasonable to adopt it, given the urgency to act.

Recommendations

1. As also urged by other Canadian researchers, e.g., by Michele Mosca and Bill Munson[9], the Canadian Communications Security Establishment (CSE) should form an advisory board to monitor and evaluate NIST's on-going post-quantum standardization effort to decide whether

[6] Michele Mosca. "Cybersecurity in an era with quantum computers: will we be ready?" In: *IEEE Security & Privacy* 16.5 (2018), pp. 38–41.

[7] For example, in 2015 it was reveal that a backdoor was introduced by the NSA in NIST's standard of a pseudo-random generator raising distrust for other NIST standards as well. (Elaine Barker and John Kelsey. "Recommendation for Random Number Generation Using Deterministic Random Bit Generators". In: *NIST Special Publication 800-90A* [2012], pp. 1–101)

[8] More information about the platforms can be found here: SUPERCOP: https://bench.cr.yp.to/supercop.html, SafeCrypto/pqclounge: https://www.safecrypto.eu/pqclounge/, pqm4: https://github.com/mupq/pqm4, and liboqs: https://github.com/open-quantum-safe/liboqs

[9] https://www.cigionline.org/articles/quantum-threat-cyber-security

to follow NIST's recommendations. This decision should, in particular, be based on the degree of transparency of the final decisions as discussed above.

2. The CSE should update (in cooperation with NIST) the Cryptographic Module Validation Program (CMVP) to require post-quantum security of cryptographic-based security systems that should protect data for medium- and long-term information lifespan as defined in the mandatory Government of Canada (GC) quantum computing threat mitigation (ITSB-127)[10].

3. The CSE should identify non-government critical infrastructures of which a lack of security pose great risk, such as in the financial market, logistics, or power plants. Similar to the ITSB-127, a mandatory risk assessment of these critical infrastructures should be implemented[11].

4. Based on the developed recommendations and standards, the Canadian Cyber Threat Exchange (CCTX) and the Canadian Center for Cyber Security (CCCS) should communicate the cyber security risk posed by quantum computers and how to mitigate it to small and large businesses and the general public.

While the second and fourth item can be approached only after a decision about which quantum-secure algorithms should be standardized, the first and third item could and should be tackled now.

[10] https://cyber.gc.ca/en/guidance/mandatory-gc-quantum-computing-threat-mitigation-itsb-127

[11] ITSB-127 only demands "communications networks, national security systems, and GC end-users who process, handle or retain GC classified information and data, or other sensitive information" to mitigate the quantum computing threat.

Data and Network Security

Quantum-secure cryptographic algorithms

Most standardized cryptographic algorithms

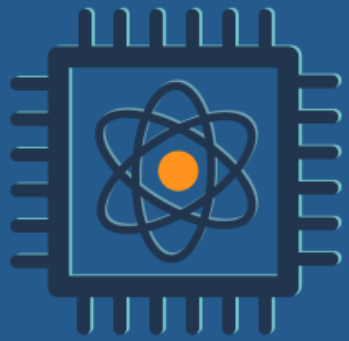

Security still guaranteed

Quantum Computer

Security NOT guaranteed

Collaborate to standardize quantum-secure algorithms as long as the foreign standardization process is trustworthy

Information gathered by Dr. Nina Bindel
Graphic created by Nate Grift

Governance of Emerging Technologies: Canadian Cyberspace Governance — or Lack Thereof?

Josh Gold, The Citizen Lab, Munk School of Global Affairs and Public Policy, University of Toronto[†]

Global Context: Governing Cyberspace

Efforts to govern the global internet, and cyberspace more broadly, are challenged by a fundamental ideological divide between countries on how they see human rights and freedoms apply to information and communication technologies (ICTs). On one side are states who frame free speech and access to information as existential threats. Led by China, Russia, and other authoritarian countries, these states have asserted sovereign control over the internet and emphasize the term "information security" (rather than "cybersecurity") when discussing threats in the digital domain. Opposing this are countries like the US, Canada, Australia and European states, to whom it is vital that the internet remain free, open, and in line with democratic values.

However, the tensions between the "cyber sovereignty" and "internet freedom" camps are not black-and-white. Liberal democracies have also increased their capabilities for control and surveillance over ICTs to varying degrees, often to address challenges such as disinformation, terrorism, or cybercrime. At the same time, cyberspace is increasingly securitized and militarized as geopolitical conflict in the cyber domain continues to intensify. Militaries and security services across the world publicly acknowledge offensive cyber capabilities—or have seen their tools leaked or stolen—as major global cyberattacks have intensified in their damage and scope.

In this context, democracies like Canada must balance real security needs and necessary defensive capabilities with the question of how to protect and promote a free, open, and peaceful cyberspace. This is difficult to do without an adequate strategy—particularly one focused internationally. Unlike many of its allies, Canada appears to lack clear, high-level strategy and policy for governing emerging technologies and the threats from them. In particular, Canada does not have a coherent cyber foreign policy,[1] and lacks transparency in its efforts to defend and promote Canadian

[†]This brief was written in the author's personal capacity and does not necessarily reflect the views of The Citizen Lab.

[1]Paul Meyer. *In 2018, will Canada finally lay out its cyber foreign policy?* 2018. URL: https://www.opencanada.org/features/2018-will-canada-finally-lay-out-its-cyber-foreign-policy/.

interests in cyberspace, raising important questions. These significant governance gaps exist despite new approaches and legal powers, a rapid and dynamic pace of innovation, and a growing rise in various kinds of digitally-enabled threats.

The State of Play: Canada

The Canadian government has moved to update policy and legislation with regard to cybersecurity. In 2018, Public Safety Canada released a National Cyber Security Strategy (NCSS), replacing the previous strategy from 2010.[2] The following year, it released the National Cyber Security Action Plan 2019-2024 (NCSAP) which stressed the need to advance Canadian interests in cyberspace internationally, while recognizing that this has not been the focus of Canadian policy to date.[3]

Regarding defence and security, the Canadian Armed Forces (CAF) announced in a 2017 updated defence policy that it would be more "assertive" in cyberspace, including the ability to conduct "active" cyber operations.[4] CAF also plans to join NATO's cyber defence centre of excellence.[5] Further, the 2019 *National Security Act* (Bill C-59) massively overhauled how the Communications Security Establishment (CSE) can lawfully operate; C-59's *CSE Act* updates the cybersecurity agency's mandate to, inter alia, include "active" and "defensive" cyber operations.[6]

Gaps, Challenges, and Forward Steps

While defence and security capabilities have rightly increased, policymaking — particularly foreign policy—has lagged behind. The lack of an international cyber strategy matters — particularly given an increasingly aggressive Canadian cyber defence posture. The Canadian government now holds that while rules and norms in cyberspace are "critical", they must be supplemented by "measures to impose costs" on hostile actors.[7] This raises questions around how Canada can support global efforts to govern the use of ICTs, while potentially using them in yet-undetermined ways as quietly agreed to by a coterie of nations.

Unlike many allies, Canada has not published its positions on how international law applies in cyberspace, despite calling at the UN for other nations to do so, and unlike its counterparts in the US and the UK, the CAF refuses to make its 2017 cyber doctrine public.[8] Further, the CAF

[2]Public Safety Canada. *National Cyber Security Strategy*. 2018. URL: https://www.publicsafety.gc.ca/cnt/rsrcs/pblctns/ntnl-cbr-scrt-strtg/index-en.aspx.

[3]Public Safety Canada. *National Cyber Security Action Plan 2019-2024*. 2019. URL: https://www.publicsafety.gc.ca/cnt/rsrcs/pblctns/ntnl-cbr-scrt-strtg-2019/index-en.aspx.

[4]National Defence. *Strong, Secure, Engaged: Canada's Defence Policy*. 2017. URL: https://www.canada.ca/content/dam/dnd-mdn/documents/reports/2018/strong-secure-engaged/canada-defence-policy-report.pdf.

[5]Josh Gold. *Canada to join NATO's cyber defence research centre*. 2019. URL: https://www.opencanada.org/features/canada-to-join-natos-cyber-defence-research-centre/.

[6]Parliament of Canada. *Bill C-59, An Act respecting national security matters, 1st session, 42 parliament*. 2019. URL: https://www.parl.ca/legisinfo/BillDetails.aspx?billId=9057418&Language=E.

[7]This was demonstrated in an October 2019 briefing note to Prime Minister Justin Trudeau—made public in January 2020. See: Jim Bronskill, "Canada ready to 'impose costs' on malicious cyber-actors, advisers tell Trudeau", 570News, January 23, 2020, https://www.570news.com/2020/01/23/canada-ready-to-impose-costs-on-malicious-cyberactors-advisers-tell-trudeau/

[8]The CAF has written a "Joint Doctrine Note, Cyber Operations", but refused to release this to the author despite

should use clear terms when referring to cyber operations; it currently labels them "active" rather than "offensive." Such opacity leads to needless ambiguity while increased transparency around Canadian foreign and defence policy positions could lead to improved signalling, thereby boosting trust, confidence, deterrence and stability.

Several other emerging security issues lack federal policy guidance. Given concerns with 5G technology, increased digital sovereignty may be needed to ensure the integrity of critical infrastructure. While it aims to be a leader in Artificial Intelligence (AI), Ottawa has not articulated clear positions on international human rights issues related to AI, including how it will support efforts to ban lethal autonomous weapons. The market for commercial spyware is widely unregulated, despite the proliferation of the sophisticated espionage and surveillance tools—which are sold to despotic regimes, including adversaries.[9] Further, given the cross-cutting nature of cybersecurity, the federal government ought to develop greater cooperation and harmonisation among various actors in the Canadian cyber ecosystem.[10]

> *"Foreign and defence policy are two sides of the same coin, and it is critically important that the two work together to build policy and strategy given the global nature of emerging technology issues."*

Challenges to Canada in the digital era may be most striking in the lack of substantive international governance considerations that the federal government has articulated to date and, where policies do exist, in their general lack of public transparency. Foreign and defence policy are two sides of the same coin,[11] and it is critically important that the two work together to build policy and strategy given the global nature of emerging technology issues.

Recommendations

While this policy brief focuses particularly on cyberspace, and how best to govern it in line with Canadian interests and values, this is but one element of national defence challenges to Canada posed by emerging technologies.

To better respond to international cybersecurity challenges to core interests, and in light of new capabilities in this domain, the Canadian government should:

---

the document having been cited in several public research papers by CAF officers enrolled in courses at the Canadian Forces College. In contrast, the US military has published its Joint Publication on cyberspace operations, and the UK Ministry of Defence published its Joint Doctrine Note on Cyber and Electromagnetic Activities—both in 2018.

[9] Bill Marczak et al. *Hide and Seek: Tracking NSO Group's Pegasus Spyware to Operations in 45 Countries*. 2018. URL: https://citizenlab.ca/2018/09/hide-and-seek-tracking-nso-groups-pegasus-spyware-to-operations-in-45-countries/.

[10] Close to twenty federal departments and agencies have an official role to play with respect to cybersecurity in Canada, in addition to relevant provincial, territorial, and municipal actors—not to mention relevant private sector and civil society actors.

[11] Sir Antony Acland and General Sir Harry Tuzo GCB OBE MC MA. "The relationship between foreign and defence policy". In: *The RUSI Journal* 128.2 (1983), pp. 3–6. DOI: 10.1080/03071848308523518. eprint: https://doi.org/10.1080/03071848308523518. URL: https://doi.org/10.1080/03071848308523518.

1. Develop and publicly articulate an international strategy for cyberspace;
2. Release its military doctrine towards cyber operations, as have the US and UK; and
3. Outline its interpretation of how international law applies in cyberspace.

Canadian Cyberspace Governance

Publish an international strategy for cyberspace

Release the military's joint doctrine on cyber operations

Articulate national positions on how international law applies in cyberspace

Information gathered by Josh Gold
Graphic created by Nate Grift

Homomorphic Encryption and Secure Outsourced Computation

Dr. Travis Morrison, University of Waterloo and Institute for Quantum Computing

Introduction

Traditional cryptography secures data while it is stored or transmitted. In general, encrypted data can not be processed until it has been decrypted. The capability of computing on encrypted data has been called the 'Holy Grail of encryption,' and has led to partnerships with 'industry players and academics to work out how homomorphic encryption could function in a Canadian setting.'.[1]

Secure Outsourced Computation

Cloud storage and computing allows an individual or organization to outsource their data management needs. If the data is privacy-sensitive, the data owner may encrypt it before transmitting it to the cloud service provider (CSP). Later, the client may wish to leverage their data in some way that requires processing the data. In order to compute meaningful statistics, for example, the data owner would have to retrieve the data from the CSP, decrypt it, and run the computations themselves, defeating the purpose of outsourcing their storage and computation needs. Such a scenario and an imagined solution were first conceived in 1978: encryption functions allowing for computation on encrypted data were dubbed 'privacy homomorphisms'.[2] In 2006, Gentry constructed the first fully homomorphic encryption (HE) scheme, which allows for computing any function on encrypted data.[3]

Homomorphic Encryption Applied

For a simple use case of HE, the client may be a hospital that uses an HE scheme to encrypt biometric patient data. The hospital then stores those encryptions in the cloud. The hospital could query the CSP for the average value of some feature, corresponding to weight or blood pressure,

[1] Catharine Tunney. *Canada's cyber intelligence agency working on 'Holy Grail' of encryption*. 2020. URL: https://www.cbc.ca/news/politics/cse-homomorphic-encryption-1.5468400.

[2] Ronald L Rivest, Len Adleman, Michael L Dertouzos, et al. "On data banks and privacy homomorphisms". In: *Foundations of secure computation* 4.11 (1978), pp. 169–180.

[3] Craig Gentry. "Fully homomorphic encryption using ideal lattices". In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 169–178.

across the patient records. The result computed by the cloud is an encryption of the average, and when the hospital receives this value, they can decrypt it to learn the average value of the feature. The CSP never learns any individual value of the weight or blood pressure across the dataset. The CSP does not learn the average weight or blood pressure either, only the encryption of these averages. Of course, the average is a simple function of a dataset. More generally, the CSP may have a model, trained by a machine learning algorithm, which can make inferences on the client's data. This model can be evaluated on the encrypted data and the encrypted inference can then be transmitted to the client.[4] Thus the confidentiality of the data is ensured: the CSP does not learn the plaintext data. Also, the CSP does not learn the true value of the inference on the data, only the encryption of the inference. Meanwhile, the client still receives meaningful insights from their data.

Going Forward with Homomorphic Encryption

It is generally impossible to update a model used for making inferences on homomorphically encrypted data. A training algorithm requires the value output by the model in order to update the model's parameters. And today, the homomorphic operations are very costly, so training a model using machine learning on encrypted data may be prohibitively expensive. This means that applying HE for privacy-preserving machine learning is best suited to a use case where there is an 'off-the-shelf' model available, like one developed by or available to the CSP. The model owner can train and update the model on other sources of data. In any case, there are other solutions to training a model in a privacy-preserving manner, such as federated learning and multiparty computation protocols. As mentioned above, there is significant overhead involved in computing on data encrypted by a homomorphic encryption protocol. Reducing this overhead is an active area of research and innovations continue to chip away at this overhead. Novel protocols increase efficiency and functionality. In addition to research and implementation, there is an ongoing effort by an open consortium consisting of experts from academia, industry, and government to standardize homomorphic encryption.[5]

> *“Homomorphic encryption can protect citizen data and privacy while allowing that data to be processed, and there is potential for applications directly related to defence.”*

Recommendations for Homomorphic Encryption and Canadian Defence

Homomorphic encryption can protect citizen data and privacy while allowing that data to be processed, and there is potential for applications directly related to defence.

1. The Department of National Defence (DND) should support collaboration between industry, academia, and other sectors of the government in this space. This is a rapidly evolving field, however. For example, from 2014 to 2019, there has been an annual competition run by iDASH

[4] Ran Gilad-Bachrach et al. “Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy”. In: *International Conference on Machine Learning*. 2016, pp. 201–210.

[5] homomorphicencryption. *Homomorphic Encryption*. 2018. URL: https://homomorphicencryption.org/.

involving one track for solving a problem in secure genome analysis using homomorphic encryption. Solving a specific problem with HE or a multiparty computation protocol may be infeasible today but it may be practical soon.

2. DND should follow the development of these technologies so they can be applied when the need arises. HE alleviates guarantees the confidentiality of private data, but a client must still trust the CSP to do the requested data processing. Homomorphic encryption allows for cloud storage and processing of data that otherwise would be prohibited due to privacy concerns.

3. It is important to analyze how HE may impact the landscape of sovereignty of information.

Homomorphic Encryption

Data confidentiality ensured

Encrypt data

Upload data

Encryption

Homomorphic encryption

Download data

Decrypt data

Process data

Encrypt processed data

Upload processed data

Process encrypted data

Transmit encrypted results

Decrypt results

Information gathered by Dr. Travis Morrison
Graphic created by Nate Grift

Human-Machine Teams in Near-Future Military Environments

Dr. Sarah Shoker, University of Waterloo

Introduction: Automation vs. Autonomy

There is no technical consensus that identifies when robots move from being automated tools to autonomous agents. However, one impactful definition suggests that in order to achieve 'teammate status,' human operators must perceive autonomous robots in the decision-making environment to be highly altruistic, benevolent, interdependent, emotive, communicative and synchronized agent teammate, rather than simply an instrumental tool.[1] Though researchers agree that the live operational environment is not yet ready for true human-machine teaming,[2] there are still important lessons that can be transferred from user experiences with automated tools to near-future operations. This is especially important as Canada begins its own path towards integrating automated tools into Canadian Armed Forces (CAF) operations, as exhibited by the decision to procure intelligent systems by companies like Calian and Plurilock. This report briefly summarizes two issues that challenge human-machine interaction today and which will continue to challenge future military environments.

Challenge 1: Algorithmic Discrimination in Human-Machine Systems

Algorithmic-driven classification and prediction can perpetuate bias against marginalized groups of people.[3] My own original research on civilian protection in post-9/11 conflict zones found that investment into intelligence, surveillance, and reconnaissance (ISR) technologies was motivated by a commitment to liberal democratic norms, including a respect for civilian immunity, but that these commitments led to paradoxical consequences. Military planners described ISR technologies as ideal tools for distinguishing combatants from civilians during insurgencies due to their sophisticated surveillance capabilities. However, human operators were still required to analyze ISR imagery. The attempt to identify plain-clothes combatants from a civilian population was challenging. Human

[1] Joseph B Lyons et al. "Viewing machines as teammates: A qualitative study". In: *2018 AAAI Spring Symposium Series*. 2018.

[2] J Christopher Brill et al. "Navigating the Advent of Human-Machine Teaming". In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 62. 1. SAGE Publications Sage CA: Los Angeles, CA. 2018, pp. 455–459.

[3] Solon Barocas and Andrew D Selbst. "Big data's disparate impact". In: *Calif. L. Rev.* 104 (2016), pp. 671–733.

operators used stereotypes related to gender, culture, and religion to distinguish between combatant and civilians, often incorrectly assuming that civilians were combatants and leading to an artificially low collateral damage count.[4] Civilian misidentification is more likely to occur in situations of irregular violent conflict, which has historically been the dominant form of political violence. However, other challenges, like correctly identifying enemy objects in visual data (e.g. identifying buildings used for housing weapons), are applicable when competing against other states.

Automated tools do not necessarily correct this problem. Classification algorithms often use some form of supervised machine learning, where a training set 'teaches' the algorithm. The training set functions as a 'textbook' crafted by human analysts who manually label the input-output data relationship in order to train the algorithm to identify similar relationships during the testing phase and then during real-world application.[5] The training data can be one of the first ways of introducing bias into the technical pipeline.[6] If the individuals responsible for labelling the training data cannot distinguish between combatants and civilians, then the machine learning algorithm will replicate the same bias. Developers who do not have first-hand experience with the operational environment may be unfamiliar with international humanitarian law and how labels contain leading language that influences operator judgements in the battlespace (e.g. fighting-age-male, adolescent, and child are frequently used by operators to make judgements.) To mitigate this problem, human operators—including software developers and individuals in the command structure—should undergo training that would help them critically examine how their own social values can influence decision-making and target selection in the operational environment.

Challenge 2: Transparency and Trust in Human-Machine Teams

AI researchers have long argued that transparency is key to establishing trust in human-machine ecosystems. Human operators must be able to audit decisions made by their AI teammates. Operators should not be overly trustworthy of AI systems, a state that is sometimes called 'artificial stupidity,' nor should an AI's design be so complicated that human operators abandon or override decisions made by AI. To avoid communication problems, other researchers have recommended that human-machine teams should use working agreements

> *"Human-machine teams should be designed so that human users can interpret, understand, and contest decisions made by the autonomous system."*

to split tasks between human and AI teammates. Human operators and autonomous agents would negotiate and allocate tasks prior to the workload reaching a "high-level".[7] This method would have human operators agree and 'buy in' to a division of labour with their AI-team members, potentially

[4] Sarah Shoker. "Algorithmic Bias and the Principle of Distinction: Towards an Audit of Lethal Autonomous Weapons Systems". In: *Digitization & Challenges to Democracy* (2019), p. 41.

[5] Google Developers. *Machine Learning Crash Course*. 2020. URL: https://developers.google.com/machine-learning/crash-course/training-and-test-sets/video-lecture.

[6] Barocas and Selbst, "Big data's disparate impact", pg 681.

[7] Robert S Gutzwiller et al. "A design pattern for working agreements in human-autonomy teaming". In: *International Conference on Applied Human Factors and Ergonomics*. Springer. 2017, pp. 12–24.

reducing communication errors and enhancing knowledge of the way AI teammates operate in adaptive environments.

Conclusion: Algorithmic Auditing

Human-machine teams should be designed so that human users can interpret, understand, and contest decisions made by the autonomous system. Autonomous technologies are especially appealing in a global environment where the limited number of human operators stall the efficiency of military operations. However, like all data-driven decision ecosystems, human-machine teaming comes with a list of serious challenges for liberal-democratic societies. AI decisions must be audited in order to mitigate biases that could:

1. undermine civilian immunity,
2. reduce human error generated by distrust, and
3. enhance communication between human and AI team-members.

Algorithmic auditing, therefore, would increase both civilian and troop protection, two norms that are very much in line with Canada's liberal democratic commitments.

Human-Machine Teaming in Near-Future Military Environments

Major Challenges to AI

Algorithmic Discrimination in Human-Machine Systems

Transparency and Trust in Human-Machine Teams

Algorithmic Auditing, enhance communication between human and A.I teams

Information gathered by Dr. Sarah Shoker
Graphic created by Nate Grift

Contributor Biographies

Authors

Dr. Nina Bindel
nlbindel@uwaterloo.ca

Nina Bindel is a post doctoral fellow at the Institute for Quantum Computing embedded in the University of Waterloo. Her field of expertise is quantum security or cryptographic algorithms. As a doctoral researcher she particularly studied the security of quantum-secure digital signatures at the TU Darmstadt, Germany. She received her PhD in 2018. In collaboration with industry and academia, she has submitted the quantum-secure signature scheme ”qTESLA” to NIST’s post-quantum standardization effort. It has been accepted as a candidate of the second round as one of nine out of 20 schemes. During the summer of 2019, she interned at Microsoft Research Redmond, USA for three months to improve the scheme, qTESLA, further.

Josh Gold
josh.gold@mail.utoronto.ca

Josh Gold (@joshgold3) is a research assistant at the University of Toronto’s Citizen Lab. His research interests include global cyberspace governance, and military and intelligence cyber operations. Josh has had prior experience at the NATO Cooperative Cyber Defence Centre of Excellence (CCDCOE), and has worked various contracts including as a consultant for the Estonian foreign ministry’s cyber policy team, and in threat intelligence at CyberCube. He is a board member of the Canadian International Council’s Toronto Branch, and of the Estonian Central Council in Canada.

Dr. Travis Morrison
travis.morrison@uwaterloo.ca

Travis is a post-doctoral fellow at the University of Waterloo with a joint appointment by the Institute for Quantum Computing and the department of Combinatorics and Optimization. Travis' research interests are in algorithmic number theory and its applications in cryptography. He received his PhD in Mathematics in 2018 from the Pennsylvania State University.

Dr. Sarah Shoker
sshoker@uwaterloo.ca

Dr. Sarah Shoker is a postdoctoral fellow in political science at the University of Waterloo where she uses empirical methods to research the security impact of emerging technologies. She is a SSHRC 2020-2022 postdoctoral fellow, along with being the winning beneficiary for the 2019-2020 University of Waterloo Trailblazer Postdoctoral Fellowship. Her work on the social impact of AI is highly requested by government policymakers. She is the founder of Glassbox, a social impact firm that trains stakeholders in government and the tech sector to identify how social values are translated into AI systems. Her Glassbox work garnered an invitation from Her Excellency the Right Honourable General Julie Payette to participate in scientific diplomacy as member of the delegation to Lithuania and Estonia. Dr.Shoker is also a member of the Government of Canada's Advisory Council on Artificial Intelligence as a member of the Public Awareness Working Group. The working group's mandate includes translating the sociotechnical impact of AI systems to the Canadian public. She was recently commissioned by Global Affairs Canada to conduct research on gender mainstreaming Canada's cybersecurity strategy; the working paper was published on the United Nations portal for the Open-ended Working Group on Digital ICTs in the context of International Security. Dr. Shoker's PhD research on military drones and autonomy in military systems was nominated for the prestigious Canadian CAGS/UMI Distinguished Dissertation Award and is currently under contract with Palgrave MacMillan, with publication scheduled for 2020.

Research Team

Kristen Csenkey
kcsenkey@balsillieschool.ca

Kristen Csenkey is the Principal Investigator of the Mobilizing Insights in Defence and Security (MINDS) Targeted Engagement Grant (TEG) to examine the future implications of dual-use and emerging technologies with military applications. Kristen is a PhD Candidate in Global Governance at the Balsillie School of International Affairs and the 2020 Canadian Global Affairs Institute (CGAI)-Women in Defence and Security (WiDS) Fellow. Her research focuses on the governance of cyber security and the management of cyber operations and innovation in Canada. She has written on Canada's contributions to international peace and security in the South China Sea, 5G telecommunication policies in Europe, and on populism in Hungary. Currently, she is a Junior Fellow with the Defence and Security Foresight (DSF) Group led by Dr. Bessma Momani (University of Waterloo) and a network member of the DSF Group's European NATO Team and the Gender Liaison. In 2016, she joined the Political Affairs team at the Embassy to Hungary, Slovenia, and Bosnia and Herzegovina in Budapest and has held other government positions since then. Kristen graduated from the University of Toronto, holds an MA in anthropology, and completed her Master of Applied Politics degree from Wilfrid Laurier University.

Chris Earle
ccearle@uwaterloo.ca

Chris is a researcher, former combat engineer for the Canadian Armed Forces, and business owner completing his Masters of Arts in Global Governance at the Balsillie School of International Affairs (BSIA). He is involved in the STEM Research Group and Conflict and Security Cluster as BSIA. As a recent SSHRC scholarship recipient, Chris is in the process of completing his Major Research Paper which applies machine learning methods to predict conflict. Chris has a Bachelors in economics with a public policy specialization from the University of Waterloo. His research applies quantitative approaches to problem solving through applied machine learning, big data analysis, and econometrics. Chris has been awarded for his honours thesis on determining incentive sources for patent filers, working with the Graham super computer cluster. He currently heads the research division of his company, Sangwa Solutions, and is a Senior Teaching Assistant for a Harvard professor-led course tailored to preparing professionals on data analytics skills to solve real world problems.

Kersty Kearney
kkearney@uwaterloo.ca

Kersty is the Research Manager for Dr. Bessma Momani's ongoing research projects, as well as the Network Manager for the Defence and Security Foresight Group, a network funded through the Department of National Defence's MINDS program. She previously worked at Interpol's Global Complex for Innovation within the 'Capacity Building and Training Directorate', where she assisted in conceptualizing and monitoring ongoing programs in Southeast Asia, including a number of efforts related to improving border security and training police forces in the region. During her undergraduate studies, she worked for the House of Commons and the Ministerial Office of the Finance Minister. Kersty obtained a Master of Human Geography from Radboud University in the Netherlands, specializing in Globalization, Migration, and Development in 2018. Her thesis sought to understand the parallels between the migration industry literature regarding migration brokers and its relevance within international sport.

Nawroos Shibli
nshibli@balsillieschool.ca

Nawroos Shibli is a Ph.D. candidate in Global Governance at the Balsillie School of International Affairs (BSIA) at the University of Waterloo where her research focuses on institutional responses to Islamophobia in Europe, with a particular focus on the European Court of Human Rights. She also holds a Master of Arts degree in Global Governance from the University of Waterloo. Previously, Nawroos served as a Senior Research Fellow at the Canadian Arab Institute and as an Editorial Assistant for Stability: International Journal of Security and Development.

Sarah Wyatt
swyatt@balsillieschool.ca

Sarah Wyatt is a former graduate from the University of Western Ontario's International Relations H.B.A. program and a current Masters in International Public Policy Candidate at the Balsillie School. Sarah is a recipient of the Global Affairs Canada Grant through her admission to the Balsillie School, and subsequently has been a participant in the Global Affairs Canada Fellowship program throughout her MIPP studies. Sarah has largely concentrated on environmental policy and security policy throughout her studies and fellowship, and insofar has become increasingly interested in the intersection between the environment and security policy due to the ever more discernible relationship between climate change and conflict. Aside from this, Sarah also maintains an interest in other topics within both the environment and security disciplines including: divestment, renewable technologies, artificial intelligence and security, and personal data protection within the digital landscape.