CANADA-UNITED KINGDOM COLLOQUIUM 2018

ARTIFICIAL INTELLIGENCE & SOCIETY

CHOICES, RISKS & OPPORTUNITIES

NOVEMBER 23-25 2018

MUNK SCHOOL OF GLOBAL AFFAIRS & PUBLIC POLICY AT THE UNIVERSITY OF TORONTO

Photos: Milan Ilnyckyj

AI & Society: Risks, Choices, and Opportunities

Rapporteur's Report

Dr. Si Yue Guo

Canada-UK Colloquium
23-25 November 2018
Munk School of Global Affairs and Public Policy
University of Toronto, Toronto, Canada

Table of Contents

MUNK SCHOOL
OF GLOBAL AFFAIRS
& PUBLIC POLICY

UNIVERSITY OF TORONTO

MASSEY COLLEGE
in the University of Toronto

Nov 23, 2018

Dear Colleagues:

The history of UK-Canada cooperation on critical policy challenges spans 150 years. This year's one hundredth anniversary of the 1918 Armistice, following a war in which Canadians and Britons fought and died together, stands as a symbol of the longevity and strength of that relationship.

The Canada-United Kingdom Colloquia has focused on many areas over the decades. By sharing research, differing approaches and respective national experiences, we have helped both countries determine and implement better economic, political, and social policies.

The immense opportunities, challenges, and risks posed by Artificial Intelligence to both countries, and indeed to the world, is fundamental to our technological, labour market, and regulatory futures. The depth and range of both delegations, and our carefully curated programme, provides an opportunity for our two countries to benefit from knowledge, research and practical experience on both sides of the Atlantic.

We are honoured to welcome our British colleagues and our Canadian delegation members to our university and city for which AI is a priority. We are eager to benefit from the insights our Colloquium will uncover and the challenges it will embrace.

The Hon. Hugh Segal Dr. Randall Hansen Dr. Peter Loewen Prof. Mel Cappe

Canada-United Kingdom Colloquium Canadian Organizing Committee Co-Chairs

CANADA

PRIME MINISTER · PREMIER MINISTRE

Welcome Letter: 2018 Canada-United Kingdom Colloquium

Canada and the United Kingdom (UK) enjoy a profound, rich, and diverse bilateral relationship. It is a relationship based on history, family, and common values. Last September, upon her first official visit to Ottawa, Prime Minister Theresa May and I committed to a Canada-UK Action Plan to enhance these relations even further through a range of new initiatives and concrete objectives.

The annual Canada-UK Colloquium contributes to our relationship by providing a multi-stakeholder platform to work together on common priorities. I am particularly pleased that the Colloquium will focus on artificial intelligence and society, a subject not only of great personal interest, but also of profound importance for Canada, the UK, and the world – now and well into the future. Canada and the UK are jointly exploring approaches to adapting to the modern economy, one where artificial intelligence, digital skills and lifelong learning will be pivotal to fostering inclusive growth.

This year's Colloquium will contribute to the dialogue to help our two nations chart future policy directions on this critical issue, which will influence modern economic, social, cultural, and political life.

It is with great pleasure that I welcome all participants to the 2018 Canada-UK Colloquium in Toronto, a city which has become a global powerhouse for artificial intelligence.

I wish you a successful event.

10 DOWNING STREET
LONDON SW1A 2AA

THE PRIME MINISTER

The UK and Canada have a strong and enduring partnership. When I visited Canada last year, Prime Minister Justin Trudeau and I expressed our joint determination to renew and strengthen the many ties of history and family that bind our two countries.

Through their annual Colloquium, for more than 45 years, the Canada-UK Council has brought together a broad range of business, academic and public sector experts to discuss a specific topic in support of our bilateral ties. The purpose is both to create networks of friendship and expertise, and to make policy proposals to both governments.

The chosen theme this year - on Rising to the Challenge of Artificial Intelligence - is particularly topical. Technology continues to advance, with Canada and the UK at the very forefront of the revolution. As I stated at the annual meeting of the World Economic Forum in Davos, it is not just in all of our interests that we should harness the power of this new technology, but fundamental to the advance of humanity. If we are to realise the enormous potential benefits of artificial intelligence we must also develop a strong ethical and regulatory framework to mitigate its risks.

I wish this year's Colloquium every success in contributing to this important work.

November 2018

The Canada-UK Colloquium

Overview

The 2018 Canada-UK Colloquium (CUKC) on "Artificial Intelligence & Society: Choices, Risks & Opportunities" took place at the Munk School of Global Affairs and Public Policy at the University of Toronto from November 23 to 25, 2018.

Established in 1971, the Canada-UK Colloquium is an annual forum aimed at facilitating in-depth discussion of critical and timely public policy issues for both countries. Every year, it brings together some fifty Canadian and British parliamentarians, senior-level public servants, academics, journalists, business leaders and young scholars, to foster practical knowledge-sharing and encourage constructive collaboration. The location of the Colloquia rotates between the two countries each year. Over the past four decades, the CUKC has been widely endorsed by the Prime Ministers from both countries, and the Foreign and Commonwealth Office in the United Kingdom and the Department of Global Affairs Canada have actively supported this initiative.

This year's colloquium focused on public policy implications of the AI revolution, both at individual and societal levels. Key topics included the development and application of Deep Learning, implications of big data and privacy and the use of AI in defence, security and policing. Special attention was also paid to science & AI, implications for the labour market and social inequality, legal liability, new regulations, and ethical considerations. Canada and the UK are already leaders in AI research and development, and the Colloquium brought together top experts to discuss ways in which both governments could mitigate the risks associated with AI and maximize benefits of new technological developments.

CUKC 2018 was co-hosted by Munk School of Global Affairs and Public Policy and Massey College at the University of Toronto.

The event was co-sponsored by the Office of Vice President, International at the University of Toronto, Microsoft, Public Safety Canada, Global Affairs Canada, McCord Travel Management, Innovation, Science and Economic Development Canada, Canadian Institute for Advanced Research, Accenture, and University College London.

Key Objectives

The objectives of the Canada-UK Colloquia have evolved over the years to focus closely on the current policy priorities of our two countries. More and more effort is being made by the British Committee of the CUKC and its Canadian partners at Queens University to ensure that they result in positive outcomes and follow-up. The two High Commissions have played a major part in assisting and encouraging this process.

Current aims are broadly as follows:

- Building up contacts between leading experts in the two countries.
- Contributing to the development of public policy in both countries.
- Identifying ways in which Britain and Canada can work together in an international context to meet their objectives in specific policy areas.
- Finding ways of achieving positive outcomes from the events and encouraging follow-up.
- Involving promising young people to ensure that the benefits of the exchanges are carried forward into the future.

AI & Society: Risks, Choices, and Opportunities

The AI Ecosystems in Canada and the UK

Canada has a vast AI ecosystem that is supported by all three sectors: government, industry, and academia, as encapsulated by Figure 1.

With three world-leading centres of excellence–Amii in Edmonton, Mila in Montreal, and the Vector Institute in Toronto–and a strong expertise in deep learning and reinforcement learning, Canada recognizes the opportunity for AI to transform a broad range of industries and improve the life of all Canadians. The Government of Canada is focused on four objectives to support AI:

- Continuing to invest in intellectual capital, in computing and data resources, and in fundamental science to support the work of researchers, while ensuring greater diversity and gender balance.
- Creating innovation superclusters and strategic innovation funds to attract business investment, as well as ensuring adoption of AI in government.
- Developing public trust and confidence in AI by examining the future of work, the ethics of AI, and marketplace frameworks through multi-stakeholder engagement.[1]
- Advancing international collaborations, which currently involve France, the United Kingdom, the G7 and G20, especially in cybersecurity.

Moving forward, Canada must look at AI holistically to build public trust and expand the AI ecosystem.

In 2017 the Canadian government announced a $125 M **Pan-Canadian AI strategy** to be implemented by the Canadian Institute for Advanced Research (CIFAR). Many key advances in AI research were enabled by CIFAR's global research programs: "AI, Robotics and Society" in 1983 and as of 2004 "Learning in Machines and Brains".

[1] Local initiatives also align with these objectives, for instance the Montreal Declaration.[1]

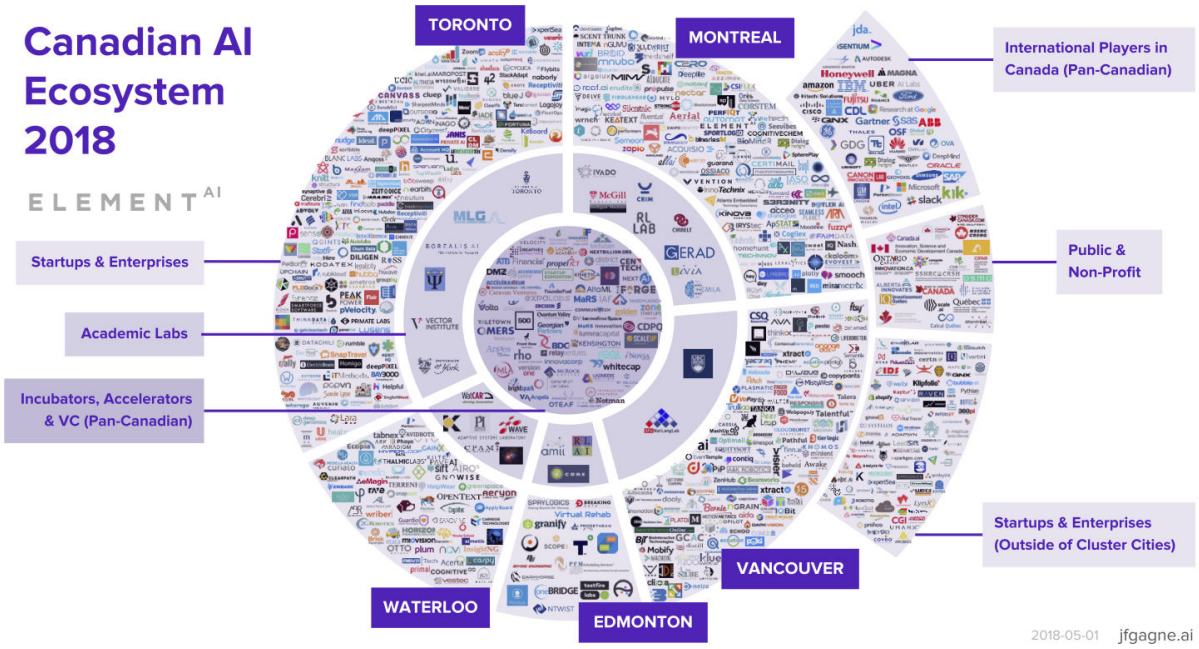

Figure 1: Canada's AI Ecosystem in 2018 [2]

The four pillars of the Pan-Canadian AI strategy, as led by CIFAR, are:

- **Centres of Excellence for AI Research and Innovation** in three major centres for AI, namely the Vector Institute, Mila, and Amii.
- The **recruitment and retention** of AI researchers and skilled graduates in Canada through initiatives like the Canada-CIFAR AI Chairs.
- Pan-Canadian activities to **support the national research community** on AI (summer schools, workshops, conferences) with equity and inclusion in mind.
- Workshops on **AI and society** to investigate the economic, ethical, and legal implications of AI and to explore long-term policy issues and opportunities raised by AI.

CIFAR has had measured success over the past two years in terms of the growing number of researchers and the significant amount of investment from multinationals, many of which created AI R&D labs in Canada.

* * *

The UK's AI ecosystem has six foundations of strength that position it alongside Canada as a global leader in artificial intelligence technologies:

- A world-class research base with a significant body of influential publications.
- A talent pipeline supported by elite universities connected by the Alan Turing Institute.
- A thriving ecosystem attracting substantial amounts of venture capitalist investment and supporting an ever-growing number of AI startups.
- Trusting, universal public institutions committed to AI-enabled innovation
- Strong appetite among local clients to adopt these new technologies.
- A regulatory approach on emerging technologies that enables innovation at pace while commanding public confidence.

To build on these strengths and maintain a world-class ecosystem, in 2018 the UK established the Office for Artificial Intelligence to deliver on the Industrial Strategy AI Data Grand Challenge "to put the UK at the forefront of the AI and data revolution". Informed by recently published independent and parliamentary reports [3, 4], its principal priorities are:

- In partnership with the Open Data Institute, improve access to data by establishing **data trusts**: frameworks to support the safe, fair, ethical and legal sharing of data to power AI innovation. Data produced through government funded research should be machine readable and open wherever possible.
- Improve supply of **skills** by attracting researchers and students, by establishing industry-funded machine learning masters, international AI fellowships, and online continuing professional development courses.
- Maximize UK **AI research** by establishing the Alan Turing Institute as the national institute for AI and data science, and making it easier for industries to license intellectual property from universities.
- Support the **adoption of AI** in the UK by working with the Government Digital Services to identify and develop the most significant opportunities to transform public services with AI technologies.

The Office for AI will work closely with the AI Council - a committee bringing expertise from industry and academia into the heart of Government - and the Centre for Data Ethics and Innovation, identifying measures needed to ensure that the development of AI and other technologies is safe, ethical and innovative.

Taken together, these institutions are developing the foundations required to increase demand and uptake which in turn creates an increase in productivity through efficiencies, but this has to be done by ensuring that the AI solutions being adopted are ethically sound, safe and work for society.

Canada-UK relationship on AI

A relationship underpinned by the 2017 Memorandum of Understanding on science and innovation signed by both our Prime Ministers and one that encompasses strong links across:

- Research: with AI a key focus of the UK-Canada Science and Technology Committee, hosted by the Alan Turing Institute in September 2018.
- Industry: with the UK-Canada Innovation Challenge with Bombardier in 2018, reinforced by our leading companies DeepMind opening its first global office in Montreal and Element AI launching its European HQ in London.
- Public Policy: sustained at the Permanent Secretary level Working Group on the Modern Economy with data-driven innovation a key focus.

The Colloquium builds on this strong foundation and gave some inspiring directions on how it could be deepened and broadened by bringing together UK and Canadian leadership in AI, including in data sharing frameworks and future research programs.

A Primer on the AI Revolution

While the industrial revolution saw machines extend humans' mechanical power, the AI revolution will see machines extending humans' cognitive power.

Most of what is termed artificial intelligence nowadays is **"weak" or "narrow" AI**. In contrast to "strong" or "general" AI, weak AI applies intelligence to specific problems rather than the entire spectrum of human abilities. It largely involves machine learning (ML) and deep learning (DL) which, due to advances in algorithm design, increased computing power, and larger volumes of data, have become easy and accessible. Examples include training continuous latent-variables, outputting grid-structured objects, and training large supervised models.

A survey of the state of existing learning techniques:

- Supervised deep learning has over the past decade reached human-level performance in many areas, including image recognition (ImageNet) and speech recognition.
- Unsupervised learning, which seeks to find "interesting" patterns in data sets, remains a difficult and possibly ill-defined problem.
- Reinforcement learning, learning by trial and error using a system of rewards and trade-offs, has enabled computers to finally beat humans in the game of Go (AlphaGo).
- Generative models that produce voice, images, natural language, and robot behaviours are beginning to teach machines to imagine.

Despite the success of machine learning and deep learning, numerous challenges remain. In particular, it is difficult to work with discrete structures, e.g. as generating large free-form discrete objects and training models with discrete latent variables, and learning from small, heterogeneous data. These arise mainly due to the absence of gradients on which to apply gradient descent optimization. Although reinforcement learning (RL) endeavours to address these problems, state-of-the-art algorithms still rely heavily on trial and error. Deep neural networks are also incapable of abstracting at multiple spatial and temporal levels, learning cause algorithms AI systems, interpreting the outcomes, communicating value, and validating and certifying AI systems.

Humans will not be completely replaced. The near-term industrial impact will be a decrease in the cost of tasks and menial work, but not a generation of new abilities. For the long term, AI researchers advocate for skepticism and caution given the amount of work needed to move from proof-of-concept algorithms to actual use. They want to avoid overhyping AI and setting the public up for disappointment. This may lead to a retreat from AI funding and research, resulting in another "AI winter".

* * *

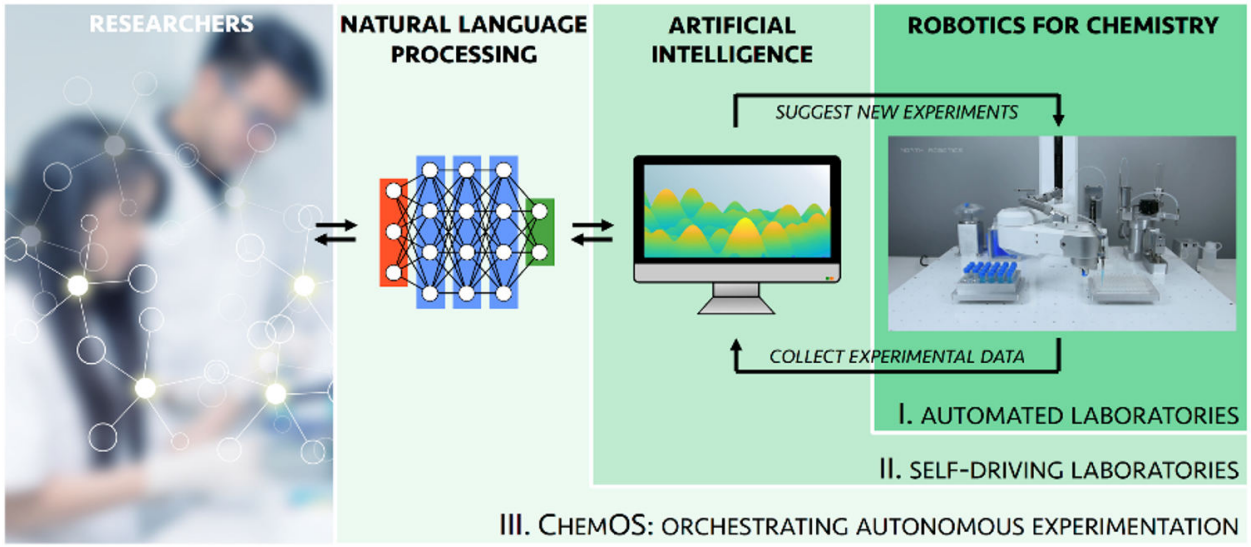

Figure 2: The closed-loop approach to discovery implemented within the ChemOS software. From Ref. [6]

Meanwhile, scientists are using machine learning, deep learning, robotics, and generative models for **materials discovery**, to **solve urgent problems** in our world like renewable energies and energy storage. With near-infinite possibilities of chemicals and materials, AI is being used to predict and screen for desired physical and chemical properties and to optimize experimental conditions.[5] The pipeline for materials discovery is a linear process in which chemical synthesis is a large bottleneck. Figure 2 shows how AI and robotics can accelerate this process and innovate global clean energy technology.

* * *

Companies and organizations are also looking to deploy AI to improve efficiencies, reduce risk, and attract and retain customers. To generate true **economic value** for everyday lives, the services they deliver must be adaptable, flexible, and personal. Achieving this requires developing and supporting four other technologies, connected devices, sensors, dynamic user interfaces, and network and cloud services, working together seamlessly in conjunction with AI.

The promise and potential impact of AI has attracted many technology entrepreneurs. Startup companies cannot afford to take on any technical risks and rely on big tech companies like Google to transform new discoveries from academia into ready-to-use software and tool kits (e.g. TensorFlow). Canada needs large AI companies to act as **"anchors"** in a growing ecosystem. These anchors, in turn, benefit from a large talent pool and from selling their support services to these startups.

Moreover, the implementation of AI must not only be accessible to developers but also to users which means online, real-time, and available on a variety of devices including mobile phones. From a software engineering perspective, the complexity and opacity of DL algorithms, e.g. neural networks, complicates white box and black box testing.

* * *

An important large-scale application of machine learning is in **information access systems**, for example web search and recommendation systems which aim to match the most relevant results to a user based on their query. These systems require a large population of users to provide feedback on the performance of the algorithm allowing it to improve over time. However, social data created by humans are inherently biased. Systems designed without considering the social implications of the algorithms were found to propagate disinformation and increase under-representation.

A potential solution to systemic bias, or algorithmic unfairness, lies in "explore and exploit systems", which through randomization allows an algorithm's performance to improve across all groups and not only the dominant one. System developers should be cautious in their experimentation, since under-represented groups tend to be over-targeted.

To address algorithm fairness, accountability, transparency, and ethics, Microsoft Research Montreal has assembled a multidisciplinary team of computer scientists, social scientists, lawyers, psychologists, etc. to form FATE (Fairness, Accountability, Transparency, and Ethics). The group develops quantitative and qualitative techniques to guide the design and application of computational systems.

Currently tech companies are struggling to measure the authoritativeness and reliability of search results and the ripple effects of propaganda, but resolving this issue is a priority.

The Politics of AI and Automation

A broad adoption of automation and AI is expected to revolutionize economies and societies, as Figure 3 shows. Its effects on politics should be considered, especially at a time when nativism and populism are rising in many countries around the world.

Parallel studies conducted in United States and Australia showed that populists and nativists in both countries are more **fearful of automation** than the average population and they are looking to governments to address their concerns. However, so far the specific **policy preferences**, e.g. reduction in immigration vs. investment in retraining, do not seem to be currently tied to their populist and nativist beliefs, although the situation is ripe for exploitation.

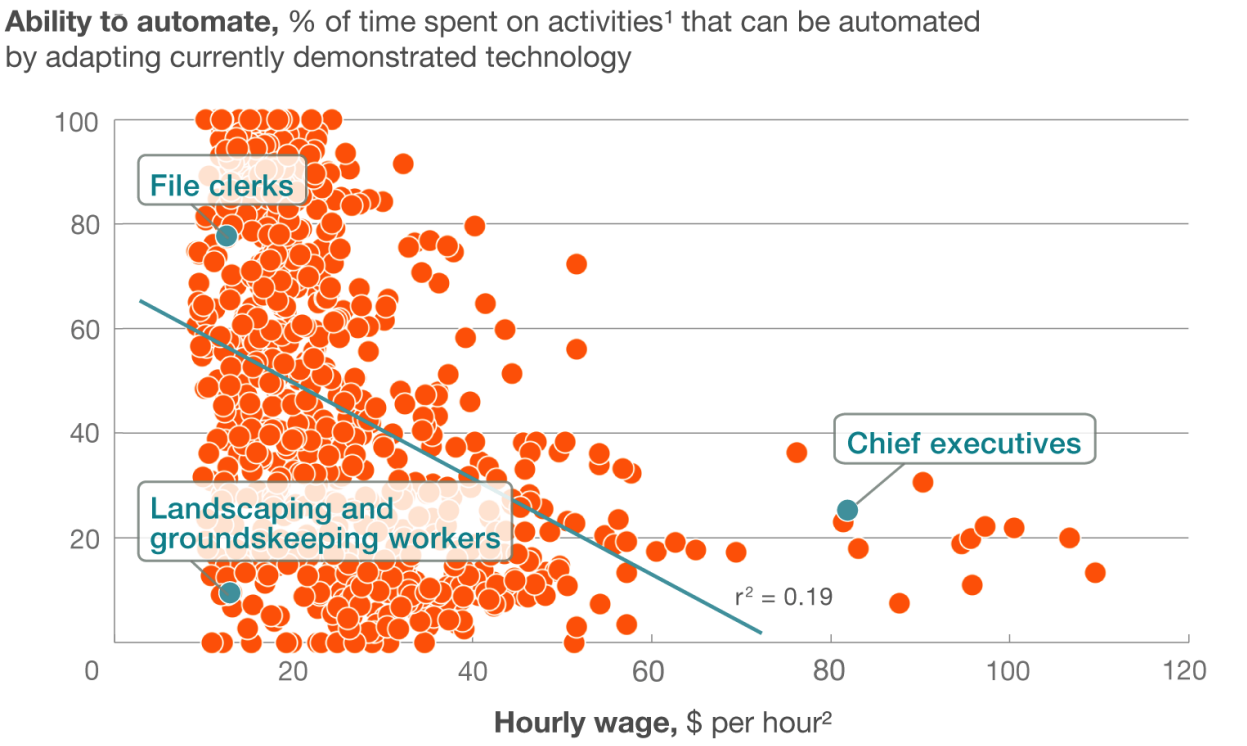

[1]Our analysis used "detailed work activities," as defined by O*NET, a program sponsored by the US Department of Labor, Employment and Training Administration.

[2]Using a linear model, we find the correlation between wages and automatability in the US economy to be significant (p-value <0.01), but with a high degree of variability ($r^2 = 0.19$).

Source: O*NET 2014 database; McKinsey analysis

McKinsey&Company

Figure 3: Comparison of wages and automation potential for US jobs. [7]

AI for Good - Ethics, Policy, and Governance

Vox Pop Labs develops AI-based **digital engagement** applications to increase political knowledge and democratic participation. These platforms operate faster and more cheaply than traditional survey methods and were found to better account for under-represented populations.[8] At a time when traditional polls are distrusted, these digital engagement tools are combating the spread of disinformation and increasing youth voter turnout. AI algorithms are also being used to analyze social media accounts for propaganda, extremist sentiment, and human rights violations.

* * *

The government of Canada is investigating the effect of emerging technologies like AI on **foreign policy and global dynamics**. While social divisions and human rights issues have always existed, AI and automation can reflect and augment societal problems.

Data ownership and privacy is another concern. As data grows in importance as an economic commodity, it becomes increasingly vulnerable to targeting and manipulation by domestic and foreign actors. Since this space is heavily controlled by private sector players, governments need to work with large tech companies.

From a geopolitical perspective, AI is creating different realities and power structures around the world. There is a global race to dominate AI research and commercialization and to set normative standards, with China and the US taking contrasting approaches.

* * *

Drastic changes in the jobs landscape are an anticipated consequence of a greater adoption of AI and automation in our economy. LinkedIn holds a large collection of employment and skills data, which is used to study **labour market issues** like unemployment and underemployment. It also helps identify and close skills gaps and match people with jobs. Given the rapid changes in demand for skills, LinkedIn is also working to help people develop the skills they need in a timely manner.

* * *

Policy makers can draw lessons from the **development of the internet**, the most successful information architecture in history that was mostly left to evolve on its own. Figure 4 shows the growth in the number of internet users since the 1990's. At the end of 2018, only half of the world's population had access to the internet. Since AI and the internet are highly integrated, there is also inequality in access to the AI around the world. When the other half of the world comes online, which approach will they take: the market forces of the US, the regulations of the EU, or the government control of China?

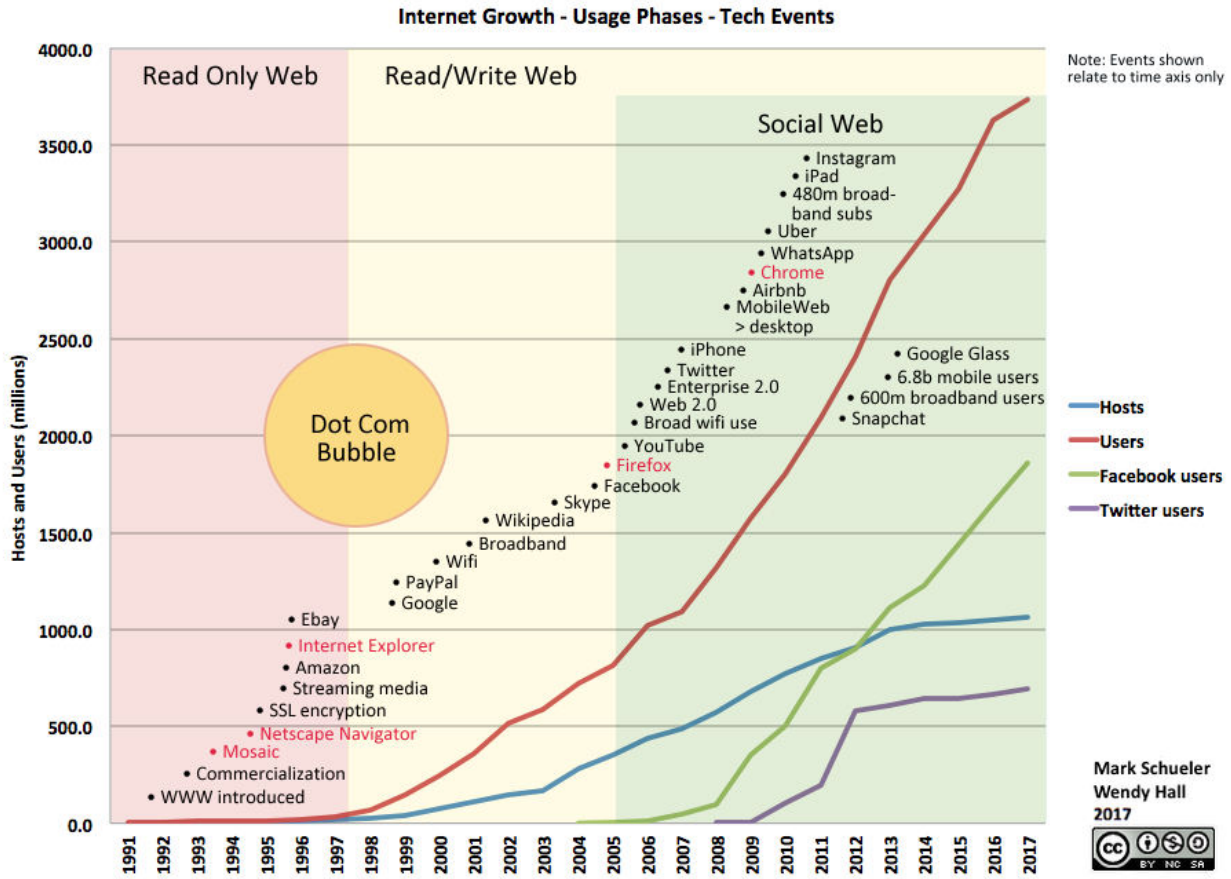

Figure 4: Internet Growth Chart. Number of users as a function of time and tech events [9]

There is insufficient action on diversity which risks importing societal biases into algorithmic biases. Of particular concern is the low percentage of women in computer science and AI, a number that is even lower in Western countries than in emerging economics.

The Legal Implications of AI - Regulations, Privacy, and Ethical Considerations

Every time a new technology is invented there is an impulse to regulate it. In reality, technology law is mostly the application of existing laws, rules, and regulations, and this is true for AI, the Internet of Things, and robots. Laws do not regulate things, they regulate people (including individuals or legal "persons" such as corporations).

The European Union's General Data Protection Regulation (GDPR) already governs the making of automated decisions by technology. There is a right not to be subjected to a

decision solely based on automated processing (e.g. profiling by machine learning with personal data) when that decision produces legal consequences or other significant effects for an individual. But if there is a human in the loop, does it fall outside of this definition? The GDPR has in place exceptions and additional safeguards like the right to human intervention and the right to challenge a decision made by a machine.[10]

On transparency, there is a requirement to provide meaningful information about the logic involved in an automated decision. But it may be difficult to comply with this requirement given opacity issues: the intrinsic "black box" nature of AI algorithms, unintelligible to the public, intentional trade secrets, and intellectual property walls (e.g. Watson's machine learning service).

On fairness, discrimination is about differentiation and is not always negative. There is significant risk of embedding prejudicial bias in both neutral discrimination and unfair and legal discrimination in training data sets. However, redacting the data to remove protected characteristics of people might undermine the potential benefits of the data set. Data minimization can collide with the objectives and the methodologies for machine learning algorithms.

Correlation is not correlation, but does it matter? It depends on the situation. That distinction needs to be made for personal healthcare and loans, but perhaps not for public health emergencies and other cases where we can act without fully understanding the causation.

There is a perception that robots are dangerous while humans are objective, fair, consistent, reasonable and unprejudiced. Regulations are often designed to protect humans from machines. In fact, humans not only exhibit prejudicial bias but human bias is often unconscious. Machine learning techniques can identify unfair bias that we didn't know existed and give us opportunities to improve the way we make decisions. GDPR allows individuals to appeal a decision made by a machine, but does not permit people to use a machine to check decisions made by other people.

A recent study deployed the Moral Machine, an experiment that asked millions of participants to resolve moral dilemmas faced by autonomous vehicles, also known as The Trolley Problem[2].[13] The outcome of this crowd-sourced exercise in moral philosophy revealed divergent cultural norms and expectations. For example, in countries where people highly respect the elderly they are more often spared from being hit by the hypothetical autonomous vehicle.

[2]The Trolley problem was devised to shed light on philosophical principles, on the theory of utilitarianism and the doctrine of double effect, initially in the context of abortion.[11, 12] It was never meant to be solved. Its purpose was to find a dilemma that would have traction on both sides. The lesson for policy makers is that they have to consider ethical design at the outset of building technologies.

The Moral Machine experiment highlights the **importance of default settings**. The devil is in the defaults. When Facebook introduced privacy settings, its default settings were set very low. Since most users neglected to changed them, the company had access to even more data than before. In the context of autonomous vehicles, regardless of what model we adopt there will be redistribution of deaths. Value choices will have to be made.

Given the great number of declarations and publications on principles for AI, we now need to move to the granular level. How do we balance principles-based and rules-based approaches? Principles work well where individual actors can be identified and risks assessed and managed in advance, subject to ex-post review. This may not be possible in a complex networked ecosystem (e.g. autonomous vehicles), where detailed rules can be automatically implemented in advance by numerous actors.

What should be done?

- Be more explicit about the ethical assumptions, assuming we agree on them and seek to engender public trust; not just "what" decisions are made and "why".
- Accept that principles-based regulation, although necessary, is not sufficient and traditional legal risk allocation models (contract, tort) most likely cannot be scaled for vast cyber-physical systems.
- Multi-stakeholder development of context-specific standards, certifications, codes of conduct of AI. More granularity is needed, i.e. extra layers but not necessarily more regulation.
- Consider collective remedial mechanisms, such as the insurance market, which is scalable, and public no-fault compensation models.

* * *

Of the numerous legal issues surrounding AI and robotics, one specific question is whether robots or AI operating independently of humans can invade our privacy.[14] Current Canadian privacy law specifies that only conscious sentient beings can invade a person's privacy. However, machines can sense, think, and act, and they are starting to outperform humans in a growing number of tasks to the point where their legal and moral status must be considered. To answer this question, three concepts were examined.

1) Human-robot interaction: Anthropomorphic design of machines has advanced to the point where users are starting to believe in the intelligence of machines (Siri, Google assistant, Pepper the robot, etc.). This trust affects how we interact with the machine.

2) The privacy doctrine: under most theories of privacy, automated processing does not harm privacy since no sentient being is involved.

3) Epistemic privacy: knowledge is the core concept to privacy. In contemporary epistemiology there is a ladder of knowledge from zero knowledge to complete certainty. In the theory of reliablism, a cognizer (human or machine) has to be capable of reliable belief formation process to enter into the realm of knowledge such that privacy can be implicated.

Since robots and AI can form beliefs about us and act on those beliefs independent of any human intervention or oversight, then yes, it can invade our privacy. The hard part is figuring out what ought to follow from this conclusion.

* * *

Concern was expressed about the use of digital evidence in legal proceedings. There is a presumption that "computers are reliable", but the test of reliability is not fully understood, despite efforts to clarify it.[15] The problem is exacerbated when software contains errors, when any single result may depend on software code stored on many distributed devices, and when no human can hope to "explain", in a court of law, how a particular AI decision was reached. It was suggested that the presumption of innocence would be eroded if there were not greater opportunity to test the reliability of computer evidence in legal proceedings.

Capitalizing on Big Data - New Frontier for Innovation

Just as we should be cognizant of the risks posed by big data and machine learning, we should also recognize the benefits for improving everyday life. Researchers at the University of Toronto have developed two projects that use smartphones to tackle mental health problems.

The first addresses social anxiety disorder and infers a patient's mental health state from information extracted from the smartphone: geographic location, samples of audio environment, physical motion of the phone, screen activity, ambient light, and SMS texts, etc. The second uses a chatbot to conduct motivational interviewing, to guide a patient to quit smoking. In both cases, the balance between privacy intrusions and patient benefits is in question. Once data is collected and human behaviour is assessed, it remains up to medical practitioners and patients how to use that information.

* * *

The Open Data Institute works with companies and governments to build an open, trustworthy data ecosystem. Data can be thought of as infrastructure, akin to roads and highways for human transportation. Just like roads and highways, society must invest it high quality, standardized data so that is reliable and useful for innovation and decision making. How do we define data ownership with regard to individuals or groups of people and also nations and the world?

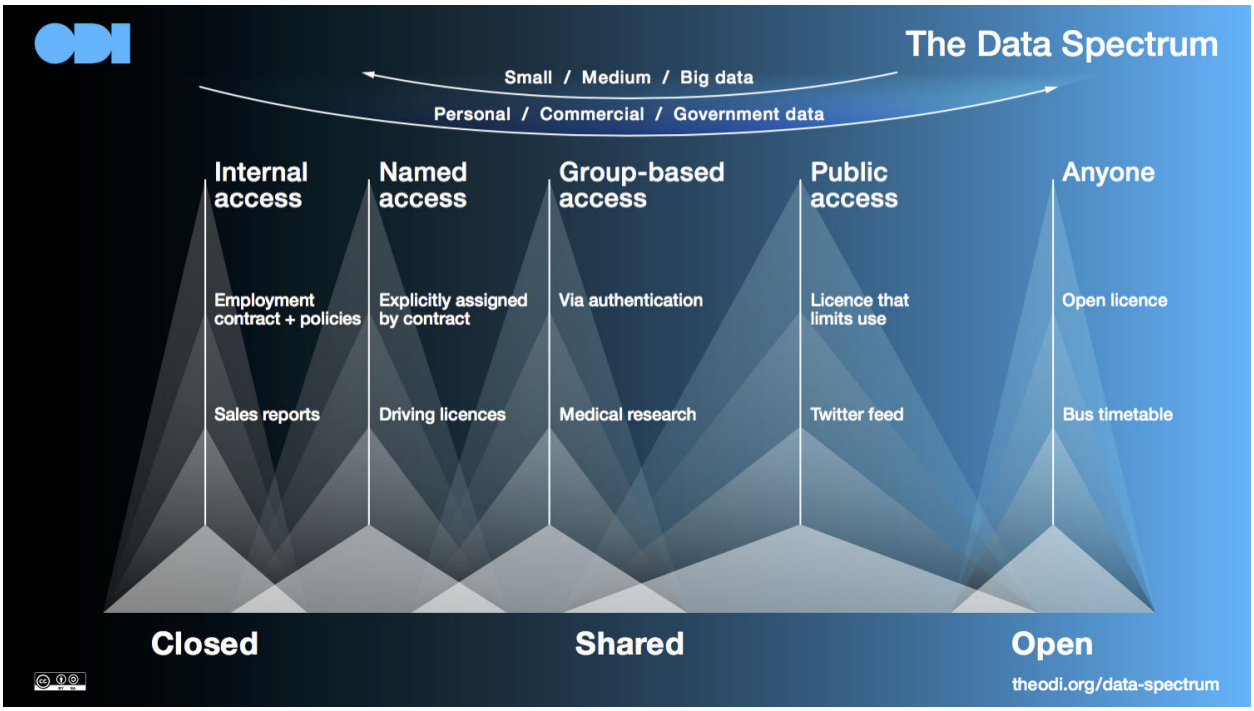

Figure 5: The Data Spectrum, from the Open Data Institute. [16]

Public buy-in and consent is essential. The discussion about data ownership and access should use language that is clear and familiar to everyone. The Open Data Institute proposes a simple classification system, shown in Figure 5. Where the data lies on the spectrum of closed to shared to open can inform us on how it should be licensed.

The integrity of review processes, such as research ethics boards, helps the public maintain trust in the system. One suggestion combines the methods of the University of Toronto research group with the concerns of the ODI: develop a chatbot to help computer scientists explore the ethical implications of their work. The results of the process can be made public. Questions remain as to how to measure the value and quality of data sets and how to incentivize companies to share data.

AI in Defence, Security, & Policing

The Canadian Navy recently conducted a strategic review and found a need to leverage data and business intelligence tools. AI is being used alongside virtual reality and 3D printing to support maintenance routines, to assist operators with monitoring, to help with recapitalization, and to support their members especially during long deployments. The goal is not to replace people but help them do their jobs better, and to improve quality of life.

* * *

The Royal United Services has examined military capabilities to identify specific areas where AI can have the greatest impact: namely, in informing, commanding, and operating. In the near-term, there are challenges regarding the amount of data, the ability to measure outcome, and the dynamic nature of the warfare. There are concerns about whether autonomous weapon systems should be allowed to make choices, especially with Russia struggling for power and influence and China acting increasingly assertive.

* * *

The British and Canadian armed forces acknowledge that in order to attract people of different technical backgrounds and skill sets, they need to adapt their recruitment and training. The rigid structure and bureaucracy of the military also risks stifling innovation. Tangentially related to AI is the major concern of cyber-warfare and cyber-espionage, especially from Russia, China, non-state actors, and criminal syndicates. These are much cheaper than traditional weapons and modes of warfare and more difficult to trace. We need to improve traceability and to engage in deterrence.

There was concern about the use of lethal autonomous systems and we should examine whether a new type of non proliferation treaty was needed. Generally, however, the impact of AI-based systems will probably be less in defence than in the wider economy and society because of the peculiarities of the sector and not least because of the limitations of data.

AI & the Work Force - Automation, Employment & Productivity

During a time of great acceleration in AI research and application by industry, the UK government is trying to manage AI use and help build toward an AI-powered future for everyone. For example, AI in the energy, healthcare, food distribution, and agriculture sectors is estimated to drive a 25% productivity boost. The nations that lead on the adoption of AI are likely be the economies that define the 21st century. Therefore, the UK needs to consider the steps to ensure they are one of those economies. Three new measures are the creation of the Office of AI, the Centre for Data Ethics and Innovation, and the AI Council. These institutions will bring the right expertise to government, to identify opportunities and challenges, and advise on policy and regulations.

One opportunity lies in the future of public service. The office of AI is working with the government digital service to identify the best areas to introduce AI and automation ahead of spending reviews for multi-year projects. The public service can help scale the workforce for AI by educating senior officials on funding for AI and by increasing efforts on diversity to have more representation in AI. The Office of AI hopes to change the perception of AI for a select few towards “AI for everyone”, a mentality which seems to be more prevalent in Asia.

We need to be open and honest with the public about the way AI and automation will change the workforce. Workers will gain the freedom to focus on creative tasks; however, many will need to undertake retraining and reskilling. Government and industry can assist in that process by creating pipelines and roadmaps for people to acquire new skills, and by working with organizations that deliver this training.

There is also an opportunity to renegotiate a new social contract, based on what we value in society. Do we want careers or contract work, and what does that mean in terms of workers rights, unions, and civil liberties?

* * *

Canadian National Railway (CN) has adopted AI to improve their operations and to help workers perform their jobs better, faster, and more safely. For example, machine vision is used for automated train and railroad inspections and maintenance. Algorithms enable them to move from a reactive to a preventative approach. When deploying technology, CN actively works with labour unions to ensure that people understand the advantages of new technology and assures workers that these measures are not intended to drive down headcount.

Breakout Sessions

Robots, AI in the Cloud, & Changing Dynamics in the Workplace

Discussion questions

- What is the correlation between robotics and job loss across various industries?
- What are the needs of employers from their employees in the new high tech workplace?
- How will AI and robotics impact extractive industries?
- Will algorithms replace management in care industries?

Summary and recommendations

- An evaluation of the sectors currently using AI found that the three most successful at implementing the technology are the financial services, retailers (understanding consumers and filling their needs), and oil and gas extraction. We can learn from the successes and failures of various sectors.
- In the changing jobs landscape we need to contrast job losses with job enrichment. We must help people to stop fearing the losses and consider what we can offer them in return. There is a potential for creating new jobs with better conditions such as shorter work weeks.
- There is concern about the quality of AI knowledge at various levels of society (government, policy makers, politicians, etc.).
- We should ensure that the two countries' strategies are not merely tinkering at a superficial level but in fact presents vision and depth.
- The question of geographic inequality (e.g. London vs the rest of the UK) is not settled: can we and do we want to make everyone equal, or do we focus on sectors of excellence in various locations?

The following recommendations were made:

Regarding employees:

1. Work with existing unions or a new type of union to focus on planning for the future.
2. Establish incentives for companies to retrain employees.
3. Establish incentives to help employees to move around.

On technology:

4. Build rigorous and common standards for AI.

5. Clarify the definitions in AI and specify the AI communitys role in explaining the technology to the public.

For the economy:

6. Encourage the larger institutions to engage with smaller companies and startups, to ensure that they receive adequate mentorship and knowledge on the transfer of skills.

7. Increase salaries and fellowships for academics working in machine learning and AI in the UK to compete with industry and other countries.

8. Empower the national institutes in the UK, like the Alan Turing institute to becomes global leaders. Learn from collaborations with the Canada institutes, Mila, Vector, and Amii.

The AI Economy: Regulations & Legal Liabilities

Discussion questions

- How do we mitigate harms whilst obtaining the benefits of AI?
- When is an ethical framework or regulation and legislation appropriate to prevent those harms?
- Can regulation be a positive for innovation?
- Do we need different rules of liability and accountability for AI?
- How can we develop a common ethical or regulatory framework internationally?

Summary and recommendations

Mitigating the harms posed by AI while maximizing its benefits requires a nuanced approach. First, we must distinguish between harms and perceived harms. Perceived harms can be minimized by building trust and trustworthiness through transparency, intelligibility, explainability, etc. We must be careful not to deem certain areas as purely harmful or purely beneficial. The trade off between individual interests and public benefit from data is a difficult problem to wrestle with, and a very germane one for organizations like the National Health Service. Lastly, ethical frameworks already exist and have converged around fairness, transparency etc. Now is the time to focus on transitioning from high level principles to implementation, assurance, and compliance.

There is a broad acceptance of the need for government to **regulate and legislate where necessary**. Some areas may not need to be covered and some areas are already covered by existing legislation on technology, human rights, anti-discrimination, etc. We need to map areas of concern where there are currently gaps. Currently facial recognition seems completely unregulated. Where appropriate, diverse and balanced public or independent oversight and auditing committees should be put in place.

Governments should be careful **not to over-regulate**, imposing a "tax" on corporations to do business. Large companies like Microsoft, Google, and Amazon, could assemble teams to work toward regulatory compliance (e.g. GDPR). However, small and medium companies may not have sufficient resources to follow complex and potentially conflicting laws from many jurisdictions. Companies welcomed GDPR, a unified system for the entire EU.

Regulation can be positive for innovation and can reduce risk for industry. Both government and industry want more clarity and discussion about what is permitted in areas where there is consequential impact on peoples' civil rights. Methods could be put in place to satisfy regulators regarding ethical decision making, e.g a regulatory sandbox

where safe conversation can take place.

Appropriate standards should be set for each layer and stage in end-to-end systems of AI, whether at input, process, output, or action, and from granular to high level. We need to consolidate the toolbox available for developing AI, including general policy, ethical design principles, regulation, incentives, shifting liabilities and insurance. Building these requires balancing transparency with commercial confidentiality, the views on which differ by country. For this to be effective there must be trust and investment in the process by all parties including the private sector, civil society, and government. In the UK, the Centre for Data Ethics and Innovation is setting the foundations.

International efforts are needed but it is more complex than the domestic situation in any one country. There is room to build on GDPR with other bilateral agreements. (The 5 eyes intelligence arrangements could complicate matters in terms of data protection adequacy with EU, especially after Brexit). We can start by mapping the systems in place and assess those that are working well within each government. We should build on common ground of Canada and UK to create bilateral consensus for AI ethical norms and legal assumptions as a first step towards potential multilateral entente.

AI & Social Inequality

Discussion questions

- What are the effects of algorithmic profiling on social cohesion and inequality?
- How do we ensure representation across all social sectors in datasets, development processes, and policymaking?
- What can we learn from the Chinese social benefit system about the improper use of AI to predict citizen's life patterns by the state?
- Should there be a set of consumer and citizen rights codified with respect to the use of AI by corporations and governments?
- Are our respective governments capable of absorbing and using all the data being collected? Are new protocols necessary? Where is Freedom of Information in all this?

Summary and recommendations

Regarding social inequality and bias, the function of machine learning and deep learning algorithms (esp. classification) is to **discriminate**. The problem arises when this discrimination leads to unfair outcomes. Hence AI models must be tailored to each application, with stipulations of what constitutes an unfair outcome and how to avoid it.

It is unlikely that we can enforce "no-go" areas for AI. However, within each application of AI, there are different levels of comfort that need to be specified.

Education of the public is key. The public need to understand what it means to consent to their data collected and shared and what the data could be used for. Governments and the private sector should use simple and clear language. Consideration could be given to how public broadcasters (BBC, CBC) could play a role in educating and engaging with the public.

If AI promises a new world, it is vital that we do not allow it to exacerbate existing social inequalities. We need to consider how to ensure an **equitable redistribution of wealth** gained from AI. Taxation could be part of the solution. The cross-border regulation and taxation of tech giants should be discussed in an international forum. Universal basic income is another possible scheme. However, in addition to income, people derive from work a sense of purpose, of self-worth, and of a place in society. The dignity of human life should be considered when governments are designing legal frameworks.

Participants noted the scale of **China**'s ambition and potential impact on the development of Artificial Intelligence technologies. Both the Governments of UK and Canada have set the ambition that future AI technologies should be developed so that they

meet societys expectations on ethics and inclusion. Participants from the research community proposed that realizing this future for global society should involve engaging a dialogue on AI with China.

Inclusion and diversity need to be considered on multiple levels. We need diversity in data collection, to ensure a well-represented data set. We also need diversity in the tech sector, in the people who are developing algorithms and who can critically assess the outputs and recognize problems. Diversity is not just about including more women and ethnic minorities, but also other under-represented groups such as the poor and older population.

1. Set clear definitions when talking about AI.
2. Enforce documentation related to data-sharing to use clear, accessible, and unambiguous language.
3. Given China a seat at the table.
4. Support research into what skills will be needed in the future workplace.
5. Explore which international forum is best suited to discuss cross-border tech giants tax liabilities.
6. Showcase success stories and role models from a diverse background, along with tangible pathways into industry.
7. Set up an independent review citizens panel to oversee the use of AI by government.
8. Study the Canadian financial regulations that allowed country to weather the 2008 financial crisis better. What can be learned from those practices that prevented economic disaster?

Full Speed Ahead - How Driverless Cars Could Transform Cities

Discussion questions

- How fast is technology advancing?
- Which comes first, smart cars or smart infrastructure?
- Are regulators keeping pace?
- What about cybersecurity and privacy?

Summary and recommendations

General Motors (GM) is committed to a future of: 1) zero emissions (electric vehicles), 2) zero crashes (autonomous vehicles), and 3) zero congestion (shared vehicles).

The discussion on driverless vehicles focused on the current state of the **technology** and how AI is integrated with existing technology. Companies have begun testing autonomous vehicles, cars, taxis, and eventually trucks and buses, in high-density urban areas. However, the cost of batteries and sensors is a major factor affecting widespread adoption.

There are questions about **public trust**, and how much more trust needs to be placed in autonomous vehicles compared to another human driver.

It appears that **regulators** are not keeping pace. The Transportation Safety Act may have to be rewritten to accommodate driverless vehicles. For instance, it currently requires that hands be on the steering wheel.

Cybersecurity is a major concern. Currently, GM vehicles experience tens of thousands of attacks on their vehicles each day.

There is much work to be done to establish proper **infrastructure** for driverless cars. Changes to be considered include decreasing speed limits, designating driverless car lanes, installing charging pad lanes, and adding sensors to the road. In major cities, driverless vehicles must be integrated with public transportation and cycling infrastructure.

Car manufacturers are looking to new **business models** that move away from solely producing cars to providing a service, such as shared taxis. Driverless vehicles that transport packages can also become shared resources for a city. These business models will likely disrupt other industries, especially the insurance industry.

We should also plan for the possibility of local or larger **monopolies**, when one company has control of the majority of cars in a city.

Conclusion

Summary and key recommendations

(1) We are at the foothills of realizing AI technologies' full potential, with some rapid advances in recent years and early adopters in pioneer sectors indicating the breadth of which AI could transform our society and economies. Currently, most of what we consider today as AI is actually a weak form involving machine learning and deep learning applied to specific tasks. As we plan for an AI-powered future, we must also prepare for strong or general AI where a single AI solution is performing many tasks and learning as it works.

(2) The AI revolution is advancing at such a speed and acceleration that it will require careful management to ensure the best outcomes for society and the economy. It is therefore incumbent upon policy makers, industry and civil society to identify the steps required set the foundations for an intelligent digital future.

(3) Though remarkable progress has emerged in recent years, the frontier of R&D in AI technologies to realize its potential remains wide and broad in areas such as scientific discovery, causality, explainability, and FAT (fairness, accountability and transparency). Therefore, governments will continue to play a vital role in should continue to funding the basic research in AI in order to ensure sustained and continual development an avert another AI winter.

(4) The gap between fundamental research and commercial application needs to be filled by applied research and software development. Companies looking to adopt AI or develop AI applications need toolkits they can use to develop and customize AI functions as well as collect and manage data. This role is currently performed by the tech giants, but can it be someone else?

(5) It is important to continue to support the other sciences and technologies that are essential for AI to succeed (sensors, batteries, dynamics user interfaces, cloud services, data management etc.).

(6) A good AI "ecosystem hub consists of start-ups and small and medium companies working alongside universities hand large tech companies that to act as "anchors". These tech companies attract talent and can take on the risk of translating academic

discoveries into useable software. Governments need to support the establishment of these AI hubs and actively remove inhibitors that prevent small businesses in these hubs from scaling-up globally.

(7) Reaping the greatest benefits of AI and managing its harms will require real, sustained multidisciplinary work as new AI solutions are developed. This includes: computer scientists, engineers, and mathematicians, working with social scientists, lawyers, policy makers, etc.

(8) We need to invest in education. For the younger generations, what kind of curriculum will best prepare them for an AI-based future? For adults, what skills do they need and how do we continually retrain them to survive and thrive in a world of AI and automation? Governments and companies, the public and private sector need to work together to develop a roadmap for education.

(9) Inclusion and diversity cannot be a consideration only at the later stages of AI implementation, it must be considered throughout the process. We need diversity in data sets and in the people of the tech sector, of gender, ethnicity, socioeconomic class, age, and any other under-represented group. That way AI solutions will reflect the needs and wants of our society.

(10) We need to clarify and use proper terminology when talking about AI. Government, industry, academics, journalists, etc. should use language that is accurate, clear, specific, consistent, and simple for the public to understand.

(11) Public engagement is very important. The public and private sectors and academia need to be open and honest with the public about the opportunities and trade-offs of adopting AI and automation. This will help allay fears of automation and AI and ameliorate the rise of populism and nativism.

(12) Part of this engagement must include how the legal frameworks related to topics such as data ownership, privacy, and liability when it comes to the interaction between intelligent AI solutions and people.

(13) Societies need to explore a new social contract as the job landscape changes under a greater adoption of AI and automation. How can we develop labour structures that will complement rather than oppose these changes? One option is to consider a shorter working week, or a greater emphasis on arts, sport and personal development.

(14) Along the lines of the EU's GDPR, international cooperation and coordinate is needed for policy and regulations. As the Cambridge Analytica incident demonstrates, any event can affect multiple countries, in this case the US, UK, and Canada. Companies also welcome standardized systems since regulatory compliance can be costly. We also need to discuss Cross-border tech giants' tax liabilities.

(15) Canada and the UK share a history, a common language, similar institutions, and a similar culture. Canada through CIFAR seems to have had a head start on its national AI institutes and attracting talent. The UK has considered in depth issues of data trusts and privacy. Both have publicly-funded healthcare systems. What can we learn from each other and how can we work together?

(16) Standardize the data architecture. Build the data, ethics, and privacy architectures to support appropriate data sharing, while balancing data protection against the danger of stifling innovation or denying potential public benefits.

(17) China is a major player in the AI space with a growing influence around the world. Despite divergent approaches to adopting AI, China must be bought into the conversation and be given offered a seat at the table.

(18) Existing inequalities in society and around the world will lead to unequal access to AI, creating a two-tier system.

(19) Now that many principles and declarations have been put forth, we need to move from the high-level ideas to the granular level. What are the different levels and approaches for each specific application?

(20) Governments should be cautious not to over-regulate, which ends up becoming a "tax" on corporations to do business. Small companies may not have sufficient resources to address complex and possibly conflicting laws from many jurisdictions. As well, some technologies are covered by existing laws, rules, regulations, and mechanisms. Of course, regulation can also be positive for innovation but there needs to be a balance.

(21) Now is the time to focus on transitioning from high level principles to implementation, assurance, and compliance. We need to consolidate the toolbox available for developing AI, including general policy, ethical design principles, regulation, incentives, shifting liabilities and insurance.

(22) In summary, AI will be a key component that drives our future. The breadth and depth of the impact to society is too great to step back and see what happens. It is a journey that we must take with our eyes open, our minds engaged and a readiness to act.

References

[1] University of Montreal. *Montreal Declaraction for a responsible development of AI*. (Dec 2018). URL: https://www.montrealdeclaration-responsibleai.com/the-declaration.

[2] Jean-François Gagné (Element AI). *Canadian AI Ecosystem*. (2018). URL: http://www.jfgagne.ai/canadian-ai-ecosystem-2018-en/.

[3] Dame Wendy Hall and Jérôme Pesenti. "Growing the artificial intelligence industry in the UK". *Independent review for the Department for Digital, Culture, Media and Sport/Department for Business, Energy and Industrial Strategy* (2017). URL: https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk.

[4] House of Lords Select Committee on Artificial Intelligence. "AI in the UK: ready, willing and able?" *UK Parliament* (2018). URL: https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/10002.htm.

[5] Rafael Gómez-Bombarelli et al. "Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules". *ACS Central Science* 4.2 (2018), pp. 268–276. URL: https://doi.org/10.1021/acscentsci.7b00572.

[6] Loïc M. Roch et al. "ChemOS: Orchestrating autonomous experimentation". *Science Robotics* 3.19 (2018). URL: http://robotics.sciencemag.org/content/3/19/eaat5559.

[7] Michael Chui, James Manyika, and Mehdi Miremadi. "Four fundamentals of workplace automation". *McKinsey Quarterly* (November 2015). URL: https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/four-fundamentals-of-workplace-automation.

[8] Wei Wang et al. "Forecasting elections with non-representative polls". *International Journal of Forecasting* 31.3 (2015), pp. 980–991. URL: https://doi.org/10.1016/j.ijforecast.2014.06.001.

[9] Mark Schueler and Wendy Hall. *Internet Growth Chart*. (2017). URL: http://growthchart.weebly.com/.

[10] Dimitra Kamarinou, Christopher Millard, and Jatinder Singh. "Machine Learning with Personal Data" (2016).

[11] Philippa Foot. "The Problem of Abortion and the Doctrine of Double Effect". *Oxford Review* 5 (1967), pp. 5–15.

[12] Judith Jarvis Thomson. "Killing, letting die, and the trolley problem". *The Monist* 59.2 (1976), pp. 204–217.

[13] Edmond Awad et al. "The Moral Machine experiment". *Nature* 563.7729 (2018), p. 59. URL: https://doi.org/10.1038/s41586-018-0637-6.

[14] Ian R Kerr. "Schrödinger's Robot: Privacy in Uncertain States" (2018).

[15] Stephen Mason and Daniel Seng. *Electronic Evidence, Fourth Edition*. Institute of Advanced Legal Studies for the SAS Humanities Digital Library, School of Advanced Study, University of London, (2017). DOI: 10.14296/517.9781911507079.

[16] The Open Data Institute. *The Data Spectrum*. (2015). URL: https://theodi.org/the-data-spectrum/.

Appendix

Program

Friday, November 23, 2018 - Briefing Day

Campbell Conference Facility, Munk School of Global Affairs & Public Policy, University of Toronto (South House, 1 Devonshire Place, Toronto)

8:30 am - 9:00 am Breakfast & Registration

9:00 am Welcome Remarks from the Canadian co-chairs

9:15 am - 10:45 am **Panel 1 - Introduction to the AI ecosystem in Canada**
(Moderator: Prof. Mel Cappe, U of T)

- "The Canadian government & AI by Dr. Nipun Vats, ISED
- "Overview of the Pan-Canadian AI Strategy" by Dr. Elissa Strome, CIFAR
- "AI, Populism and Governance" by Dr. Peter J. Loewen, U of T
- "AI & cyber security by Dr. Fenton Ho, Public Safety Canada

10:45 am - 11:00 am Coffee & Tea break

11:00 am - 12:30 pm **Panel 2 - Delving "Deep": AI Explained**
(Moderator: Dr. Tom Axworthy, InterAction Council)

- "Machine Learning 101" by Dr. David Duvenaud, Vector Institute
- "Practical Applications of AI" by Dr. Alán Aspuru-Guzik, U of T
- "AI in Our Every Day Lives" by Jodie Wallis, Accenture
- "AI & Algo Anywhere: A Case Study by Adam Gravitis

12:30 pm - 2:00 pm Lunch and **keynote address by Fernando Diaz, Microsoft**

2:00 pm - 4:00 pm **Panel 3: AI for Good – Ethics, Policy and Governance**
(Moderator: Ms. Sonya Thissen, Global Affairs Canada)

- "AI, big data & electoral process" by Cliff van der Linden, Vote Compass
- "Responding to automation: predicting changing skills requirements and closing skills gaps using LinkedIn data & learning content" by Jake Hirsch-Allen, LinkedIn
- "AI and foreign policy: human rights and geopolitical implications by Tara Denham, Global Affairs Canada

6:00 pm Networking reception at Massey College Junior Common Room

7:00 pm Dinner at Massey College Upper Library with **keynote address by Scott Daniels** (CN)

Saturday, November 24, 2018 - Colloquium Day 1

Presiding Chair: Dr. Alan Bernstein, CIFAR

8:30 am Breakfast & Registration

9:00 am Welcome Remarks from Dr. Randall Hansen, U of T

9:15 am- 10:30 am **Session 1: Debunking AI & Deep Learning - A Primer on the AI Revolution**
CA: Prof. Doina Precup, McGill University
UK: Prof. Dame Wendy Hall DBE, Computer Science, University of Southampton

10:30 am - 10:45 am Coffee & Tea break

10:45 am - 12:00 pm **Session 2: The Legal Implications of AI - Regulations, Privacy and Ethical Considerations**
CA: Prof. Ian Kerr (University of Ottawa)
UK: Prof. Christopher Millard (Queen Mary, University of London)

12:00 pm - 2:00 pm Networking Luncheon at the Upper Library, Massey

2:30 pm - 4:00 pm **Session 3: Capitalizing Big Data - New frontier for innovation**
CA: Prof. Jonathan Rose, University of Toronto
UK: Dr Jeni Tennison, OBE, CEO, Open Data Institute

4:00 pm - 5:30 pm **Session 4: Break-out sessions**

BoG1: Robots, AI in the Cloud & Changing Dynamics of the Workplace
CA: Dr. Peter Warrian, Innovation Policy Lab, Munk School
UK: Dr. Rannia Leontaridi OBE, Director, Business Growth, BEIS

BoG2: The AI economy: regulations & legal liabilities
CA: Ms. Andree Gagnon, Microsoft Canada
UK: Lord Clement-Jones CBE, Chair, House of Lords Select Committee on AI

BoG 3: AI & Social inequality
CA: The Hon. Hugh D. Segal, Massey College
UK: Dr Matthew Fenech, AI Policy Consultant, Future Advocacy

BoG4: "Full speed ahead - how the driverless cars could transform cities"
CA: David Paterson, GM Canada
UK: Professor Mark Andrew Girolami FRSE, Alan Turing Institute

6:00 pm Networking Reception & Dinner at Hazelton Hotel, Yorkville

Sunday, November 25, 2018 - Colloquium Day 2

8:30 am Breakfast

9:00 am - 10:15 am **Session 5: AI in Defence, Security & Policing**
CA: V. Adm Ron Lloyd (Commander of the Royal Canadian Navy)
UK: Trevor Taylor (Royal United Services Institute for Defence & Security Studies)

10:15 am - 11:15 am **Session 6: AI & The Work Force – Automation, employment & productivity**
CA: Scott Daniels (CN)
UK: Sana Khareghani (Deputy Director, Head of the Office for AI)

11:15 am - 12:30 pm **Session 7: Reports from Break-out sessions**

12:30 pm Shuttle to the Hotel & Luncheon at Yorkville Room, Hazelton Hotel

1:30 pm - 2:00 pm **Session 8: Summary from the Rapporteur**

2:00 pm - 2:10 pm **Closing Remarks from Dr. Alan Bernstein**

2:10 pm - 2:20 pm **Concluding remarks from the Organizing Committee**

2:20 pm - 2:30 pm Group Photo and Adjournment

Participants

Canadian Delegates

Jake-Hirsch Allen, *LinkedIn Canada*
Prof. Alán Aspuru-Guzik, *University of Toronto*
Dr. Tom Axworthy, *Massey College*
Dr. Alan Bernstein, *Canadian Institute for Advanced Research (CIFAR)*
Dr. Mel Cappe, *University of Toonto*
Sharly Chan*, *Massey College*
Scott Daniels, *Canadian National Railway (CN)*
Tara Denham, *Global Affairs Canada*
Dr. David Duvenud, *Vector Institute*
Drew Fagan, *University of Toronto*
Everett Findlay, *Deloitte Canada*
Rebecca Finlay, *CIFAR*
Andrée Gagnon, *Microsoft Canada*
Lieutenant(N) Adam Gravitis, *Royal Canadian Navy*
Dr. Si Yue Guo, *University of Toronto*
Dr. Randall Hansen, *Munk School of Global Affairs & Public Policy, University of Toronto*
Dr. Fenton Ho, *Public Safety Canada*
Dr. Ian Kerr, *University of Ottawa*
Dr. Peter J. Loewen, *Munk School of Global Affairs & Public Policy, University of Toronto*
Vice Admiral Ron Lloyd, *Royal Canadian Navy*
Dr. Tina J. Park, *Canadian Cenre for the Responsibility to Protect*
David Paterson, *GM Canada*
Dr. Doina Precup, *McGill University*
Prof. Jonathan Rose, *University of Toronto*
The Hon. Hugh D. Segal, *Massey College*
Dr. Elissa Strome, *CIFAR*
Sonya Thissen, *Global Affairs Canada*
Dr. Nipun Vats, *Innovation, Science and Economic Development Canada*
Dr. Cliff van der Linden, *Vox Pop Labs*
Jodie Wallis, *Accenture Canada*
Dr. Peter Warrian, *University of Toronto*

Canadian Observers

Jonathan Abecassis, *CN*
Marlene Floyd, *Microsoft*

UK Delegates

Dr. Natalie Banner, *Wellcome Trust*
Dr. David Barber, *AI Centre, UCL*
Anthony Cary, *Chairman, Canada-UK Council*
Dr. Mandy Chessel, *IBM*
Lord Tim Clement-Jones, *Chair, House of Lords Select Committee on AI*
Sue Daley, *TechUK*
Dr. Kate Devlin, *Kings College London*
Dr. Matt Fenech, *Future Advocacy*
Prof. Luciano Floridi, *Oxford Internet Institute*
Prof. Mark Girolami, *Alan Turing Institute*
Prof. Dame Wendy Hall, *University of Southampton*
Dr. Julian Huppert, *Cambridge University*
Sana Khareghani, *Office for AI, DCMS - BEIS*
Dr. Rannia Leontaridi, *Office for Artificial Intelligence, UK Government*
Fernando Lucini, *Accenture*
Nicolas Maclean, *Council Member, Canada-UK Council*
Stephen Mason, *Barrister*
Rob McCargow, *PwC*
Prof. Christopher Millard, *Queen Mary University of London*
Prof. Jim Norton, *Adviser, Canada-UK Council*
Anneke Schmider, *Information Commissioners Office*
Jeni Tennison, *Open Data Institute*
John Thornhill, *Financial Times*
Prof. Trevor Taylor, *Royal United Services Institute*
Philippa Westbury, *Royal Academy of Engineering*
Prof. Michael Wooldridge, *University of Oxford*

UK Observers

Dr. Matias Acosta*, *University of Cambridge*
Dr. Saif Ahmad*, *University of Cambridge*
Angus Barry, *Government Digital Service*
Prof. Jonathan Heeney, *Cambridge University*
Gale Jenkinson, *Executive Secretary, Canada-UK Council*
Kevin McGurgan, *British Consul-General, Toronto*
William Swords, *Vice President and Chair of their International Trade Forum*
Edward Teather, *Department for Digital, Culture, Media and Sport*
Dr. Karina Vold*, *University of Cambridge*
Dr. Nikki Wekman*, *University of Cambridge*

* Notetakers

About the Author

Dr. Si Yue Guo is a postdoctoral researcher in physical chemistry in Prof. Alán Aspuru-Guzik's group at the University of Toronto where she is developing a new paradigm for computation using chemical reactions. An alumna Junior Fellow of Massey College, she has a Joint Honours BSc in Physics and Mathematics from McGill University and a PhD in Physical Chemistry from the University of Toronto, where she studied the reaction dynamics of molecules at crystal surfaces in Prof. John Polanyi's lab. Dr. Guo is interested in using computational techniques including artificial intelligence to study chemical reactions.


Adam Gravitis
Jodie Wallis
Dr. Alan Aspuru-Guzik
Dr. David Duvenhaud

2018 CANADA-U.K. COLLOQUIUM

CO-HOSTS

munk school
OF GLOBAL AFFAIRS & PUBLIC POLICY

UNIVERSITY OF TORONTO

MASSEY COLLEGE

Canada-United Kingdom Council

PLATINUM SPONSORS

Microsoft

CN

GOLD SPONSORS

accenture

CIFAR

Global Affairs Canada
Affaires mondiales Canada

Innovation, Science and Economic Development Canada
Innovation, Sciences et Développement économique Canada

McCORD Travel Management

UNIVERSITY OF TORONTO

Public Safety Canada
Sécurité publique Canada

UCL