

The AI Shift:

Implications for policymakers

April 2018



AUTHORS



SARAH VILLENEUVE

Policy Assistant

Sarah conducts research under the AI + Public Policy and Innovation + Inclusive Growth work streams. Motivated by the potential benefits emerging technology offers to governments and civil society and concerns of fairness, accountability and transparency, she seeks to contribute to critical conversations on the development and adoption of technology in public life. She is a contributor to the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems and IEEE Standards Association. She has previously conducted research on algorithmic discrimination, smart-city marginalization, and predictive analytics for governance. Sarah holds a MSc in Data and Society from the London School of Economics and Political Science and a BA in Politics and International Relations from Royal Holloway, University of London.

sarah.villeneuve@ryerson.ca
[@sarvilleneuve](https://twitter.com/sarvilleneuve)



NISA MALLI

Senior Policy Analyst

Nisa Malli leads research for the Brookfield Institute on the public policy implications of artificial intelligence and creating more inclusive innovation-driven economies. Previously she has worked on emerging technology, digital services, and public sector innovation for federal and municipal governments and managed a nonprofit digital literacy program for seniors and newcomers. Nisa holds an MA in Public and International Affairs from the University of Ottawa, where she researched open government and digital citizen engagement and a BFA in Writing from the University of Victoria.

nisa.malli@ryerson.ca
[@nisamalli](https://twitter.com/nisamalli)


CONTRIBUTORS

Brandon Chu, Photographer
Jesse Darling, Senior Projects Designer
Melinda Jacobs, Project Advisor, AI + Public Policy
Tara O'Neil, Facilitator
Michelle Park, Projects Officer
Heather Russek, Acting Director, Policy Innovation Platform

The Brookfield Institute for Innovation + Entrepreneurship (BII+E) is a new, independent and nonpartisan institute, housed within Ryerson University, that is dedicated to making Canada the best country in the world to be an innovator or an entrepreneur.

BII+E supports this mission in three ways: insightful research and analysis; testing, piloting and prototyping projects; and thoughtful policy innovation approaches. These three disciplines/ areas of focus reinforce the others, and are supported by a culture of collaboration, community engagement and impactful partnerships.

For more information, visit brookfieldinstitute.ca.

 [@BrookfieldIIE](https://www.facebook.com/BrookfieldIIE)

 [/BrookfieldIIE](https://twitter.com/BrookfieldIIE)

 [The Brookfield Institute for Innovation + Entrepreneurship](https://www.linkedin.com/company/brookfield-institute-for-innovation-entrepreneurship)



Pictured (left to right): Dr. Taylor Owen, Dr. Ben Scott, "Special Topic: Digital Deceit"

TABLE OF CONTENTS

Introduction	5	Appendix 1: Agenda	17
Event Design	6	Appendix 2: Case Study Discussion Prompts	19
Agenda	6	Appendix 3: Case Study Cards	20
Facilitation structure	7	Appendix 4: Opportunities and Challenges Map	22
Case studies	8	Appendix 5: Participant Evaluations	23
Understanding the Shift	9		
Technical capabilities	9		
Commercial applications	10		
Government response	11		
Facilitated Discussions	13		
Values	13		
Bias and discrimination	13		
Data ownership and control	14		
Transparency and open source algorithms	14		
Privacy	14		
Government regulation and consumer protection	15		
Next Steps	16		

INTRODUCTION

Artificial Intelligence (AI) has benefited from a number of recent technological advances, from increases in processing power to decreases in battery costs, and the explosion of available data. Beyond technical research and commercial opportunities, this transformative technology has the potential to fundamentally alter our society and how our public institutions operate.

Despite the scale and magnitude of the impact, conversations about AI remain fragmented and governments face challenges connecting with industry and academia on the edge of technological advances, in recruiting expertise, and in identifying and procuring solutions that could significantly improve government decision-making, operations, and service delivery.

On March 23rd, 2018, the [Brookfield Institute for Innovation + Entrepreneurship](#) (BII+E), with the Ontario government's Policy Innovation Hub, hosted a one-day conference, [AI + Public Policy: Understanding the shift](#).

This event was among the first of its kind in Canada, designed to develop a shared understanding of the core technical concepts and historical context of AI among multi-sectoral participants, and to encourage deliberation on the cross-cutting challenges and public policy implications of this evolving technology.

This dynamic conversation among policymakers, technologists, entrepreneurs, and academics was rooted in the terms and concepts presented in [Intro to AI for Policymakers: Understanding the shift](#), which was provided to participants as background reading in advance of the event.

The aim of the event was to:

- + Provide policymakers with a direct line of sight into the AI sector: myths, hype, the evolving state of technological advances, and potential applications.
- + Identify policy areas for further exploration and areas with multi-jurisdictional implications.
- + Strengthen connections between policymakers and AI experts, and between the public, private, and academic sectors.

This report:

- + Provides a summary of the insights from speakers and participants throughout the event.
- + Explores the challenges and opportunities that AI poses to society.
- + Dives deeper into the public policy implications of AI and the role of government in this evolving technical landscape.

One of the key conclusions from the day was the need for deliberate conversation among policymakers, technologists, social scientists and the broader communities that will be impacted by a shift toward a prediction-centred society. In the coming months, BII+E will continue to create space for these cross-sector conversations and collaborations, through partnerships and our own research.

EVENT DESIGN

AGENDA

This event was structured to build a shared understanding of core technical concepts and AI history through keynotes, panels, and lightning talks, and to encourage deliberation between policymakers, industry, and academic experts. Throughout the day, participants were encouraged to use [Sli.do](#), an interactive, web-based tool that enables real-time Q+A and audience polling. Speakers included academics and cutting-edge PhD researchers, industry leaders, startup founders, policymakers, and technology legal experts. (See Appendix 1 for full agenda and speaker list.)

The morning panels and keynotes provided foundational knowledge, equipping policymakers with an understanding of the historical emergence, technological components, current capabilities, and applications of AI, along with existing challenges for government.

The afternoon transitioned from level-setting presentations to collaborative workshops on the opportunities and challenges posed by AI, and presented real-world applications to prompt

discussions on policy responses and government's role. At the end of the day, 50% of participants felt optimistic about the future of AI.

“The AI Shift was one of the richest conferences on the topic of AI and public policy that I have been to. Personally, I feel like it would have benefited from a second day where we had the opportunity to dive deeply into some of the issues that were raised, because there was a lot of ground to cover in a short period of time.”

- Michael Karlin, Senior Advisor, Treasury Board of Canada Secretariat



Pictured: Michael Karlin, “So What? Implications of AI for Policymakers”

FACILITATION STRUCTURE

The first workshop centred around six cross-cutting challenges posed by AI, including the issues of:

- + Bias
- + Ethics
- + Privacy
- + Safety
- + Explainability
- + Accountability

The session opened with expert lightning talks on each topic and participants were invited to pick two of the six challenge areas that they were most interested in learning more about.

Six stations were set up around the event space, each staffed by a member of the BII+E facilitation team, with additional subject matter expertise provided by the lightning talk experts. Participants were guided through a mapping exercise, identifying the opportunities and challenges each issue posed to government and/or the private

sector. The consolidated opportunities and challenges map created by participants is provided in Appendix 4.

The second workshop focused on real-world case studies, grounded in realistic, probable (and in some cases, current) applications of the technology. (See Appendix 3 for details on case studies.) Participants were provided with 10 case studies and each table was invited to review the case studies and select one to focus on. Case studies were developed by BII+E to reflect a broad range of technological innovations, potential applications, and policy areas.

“In this period of hype, uncertainty, and concern, the AI + Public Policy conference this past Friday was a welcome breath of fresh air.”

- Tim Dutton, Master of Global Affairs Candidate at the University of Toronto




Pictured: Patryce Bowling, Facilitator, “Decoding AI: Impacts for Policymakers”

CASE STUDIES

- + **Automated Hiring** and the use of computer vision, employee data, and predictive analytics to select candidates and evaluate current staff.
- + The use of **Smart Home Technology** and data collection by private development companies and building management firms, including smart security cameras, doorbells with voice and facial recognition, and voice-activated home systems.
- + **Autonomous Vehicles** and the question of accountability in the case of an accident.
- + **Chatbots** as a tool for telehealth service delivery and diagnostics.
- + **Predictive Life Insurance** that uses data from fitness wearables and apps to determine insurance eligibility and premiums.
- + **Adaptive Education** and the use of AI assistants in K-12 classrooms and real-time data collection to evaluate student progress and stream students.
- + The use of **Predictive Legal Analytics** to analyze historical behaviour and judiciary decision-making and advise lawyers on how to best approach litigating or defending cases.
- + Using data on youth in care and predictive algorithms to **Detect At-Risk Youth** and trigger interventions from social services or police.

- + Using crowdsourced data, official reports, and prioritization algorithms to target **Road Maintenance**.
- + **Employee Tracking** using smart bands and localized sensors to log and analyze activity.

Discussion prompts (Appendix 2) invited participants to consider how AI could help, which policy domains would be impacted, who might be made vulnerable in the chosen scenario, cross-cutting issues (e.g., bias, safety, etc.), and considerations for government. Out of the six roundtable discussions, three selected Automated Hiring, two selected Predictive Life Insurance, and one chose Adaptive Education. The full case studies are provided in Appendix 3.



UNDERSTANDING THE SHIFT

The conversation that this event and invited speakers sought to unpack reflects shifts in technological capabilities, commercial applications and how government should respond.

- + **Technical capabilities** - The ability to harness data via algorithmic models, in combination with increases in computational speed, has enabled AI take on increasingly complex tasks. In her opening remarks, Dr. Kathryn Hume, VP of Product and Strategy at integrate.ai, described the shift from descriptive to predictive analytics, from understanding the past to building future models.
- + **Commercial applications** - As Dr. Kathryn Hume described, improvements in software capabilities, coupled with the reduced cost of hardware, availability of data, and the emergence of an internet-based open-source community, have moved AI from niche academic projects to a robust industry. Prediction has become embedded within products and used for processes that have traditionally been outside of the domain of machine intelligence. The cost of computation has continued to decrease, giving more companies access to the tools of prediction.
- + **Government response** - These shifts are impacting governments and their role as regulators, users, and collaborators. Faced with quickly evolving technology and industry landscapes, governments have an imperative to be both proactive and deliberate in their planning to adopt, implement, and regulate emergent technologies. The scale and speed of AI technological advances and commercial applications necessitates collaborative, cross-disciplinary discussions among policymakers, technologists, social

scientists, and members from the broader community that will be impacted by the increasing use of this technology.

These trends provoke a profound rethinking of the foundational processes by which we mark our social, economic, and political progress.

TECHNICAL CAPABILITIES

Theories and models of AI predate the availability of hardware, data, and scalable software to support their implementation. However, as data and computational speed have grown and the field has evolved, intelligent systems are increasingly able to forecast, recommend, and make calculated decisions, shifting what have traditionally been human-centric tasks into the domain of AI.

Both Dr. Graham Taylor, Associate Professor, School of Engineering at the University of Guelph, and Dr. Kosta Derpanis, Associate Professor, Department of Computer Science at Ryerson University, pointed to the ability of AI to perform certain tasks more accurately and efficiently than humans. Machine intelligence excels at identifying and sorting images at a speed and accuracy that is prohibitive to human intelligence.

In addition to augmenting human capability, some algorithms are able to extend the range of available tasks, such as solving complex equations with large volumes of data. As Dr. Graham Taylor noted in his remarks, “When you hand over the design of things like software to an algorithm, and it is viewed as essentially a search through a very high dimensional space of possible programs, these machines can build software that’s much more effective than software that humans can write and it can also solve problems that are just out of reach of humans.”

Yet tasks such as moral reasoning and intuition, which come instinctually to humans, have

remained out of reach for automation. In fact, these tasks were featured by speakers as important inputs for successful AI development. Human judgement and meaning-making was recognized as an integral part of standardizing data, training algorithms, and interpreting outputs.

In his talk on computer vision, Dr. Kosta Derpanis noted that humans are needed to standardize and label training data that is used to teach computer vision and image recognition systems. Humans effectively describe and add meaning to what would otherwise be uninterpretable data for computer systems. Dr. Graham Taylor stressed that continued advancements in computer performance, coupled with the increase of available data, will lead to even greater technical capacity and the development of new abilities and applications of AI. The more AI is capable of, the more impact it will have on the structure and value of human tasks.

Even in the weeks leading up to and shortly following the event, the BII+E team noted some new and novel applications of AI in the media, demonstrating the speed at which technical capabilities are improving, as well as challenges in ensuring policy is up-to-date with current and future applications. These included:

- + [An automated system capable of detecting diabetic eye disease](#)
- + [Spotting wildlife poachers in real-time using computer vision, deep learning, and infrared imaging](#)
- + [Predicting demand for ride-hailing services](#)

COMMERCIAL APPLICATIONS

The shift in commercial applications of AI closely mirrors advances in technical capabilities. As Dr. Kathryn Hume emphasized, improvements in software capabilities, coupled with the reduced cost of hardware, availability of data, and the

emergence of an internet-based open-source community have moved AI from niche academic projects to a robust industry.

Prediction has become embedded within products and used for processes that have traditionally been outside of the domain of machine intelligence. The cost of computation has continued to decrease, giving more companies access to the tools of prediction. Yet, as highlighted by Dr. Avi Goldfarb, Ellison Professor of Marketing, Rotman School of Management at the University of Toronto, the cost of its complements - namely input data, human judgement, and actions taken based on derived insights - have risen.

“When the price of coffee falls, we buy more cream and sugar. And so the thing you need to be asking yourself [is], what are the cream and sugar for prediction? As the price of prediction falls, what becomes more valuable?”

– Dr. Avi Goldfarb, Ellison Professor of Marketing,
Rotman School of Management at the
University of Toronto

Advancements in predictive techniques have transformed products and services from generalistic to hyperpersonal. Commercial applications have moved from a one-size-fits-most offering to an individualized marketplace of services, products, and recommendations that adapt to individual use and needs. In some cases, prediction has been used in ways that could be detrimental to society and public interest.

Dr. Ben Scott, Senior Advisor, Open Technology Institute at the New America Foundation, highlighted how prediction has been used to disrupt democratic processes through micro-targeting online political advertisements. In his remarks, he challenged the notion that greater information availability and technical capabilities

inherently benefit society, emphasizing, “We have too long believed that technology and its progress is a net good.”

GOVERNMENT RESPONSE

Government as a regulator

Speakers illustrated the necessity for government to assume a proactive regulatory role by describing both the speed at which technological capabilities are developed and adopted, and the scale at which AI is being applied in public life. In his opening remarks, Jay Porter, Director, Policy Innovation Hub, talked about the role of government in protecting the public good and society at large from the negative impacts of disruptive technology, while serving as a demand driver and supporting the innovation revolution and growth of Ontario firms.

Doing so requires policymakers to better understand the potential impacts AI can have on society. Lex Gill, Research Fellow at Citizen Lab; Abhishek Gupta, AI ethics researcher at Concordia’s District 3; and Alexandre Piché, doctoral student at the Montreal Institute for Learning Algorithms, encouraged policymakers to have more forthright conversations about issues of bias, ethics, privacy, and safety, and AI’s potential to exacerbate existing systemic shortcomings. In her remarks, Lex Gill urged attendees to move beyond individual consent models for privacy, introducing the concept of privacy as a public good and core component of democracy.

Government as a user

Alongside the conversation on government regulation, participants and speakers discussed the potential benefits of applying AI in the delivery of government service and internal decision making. Improvements in capabilities and the development of applications that employ natural language processing (NLP) and computer vision provide opportunities for government to improve services

and internal efficiency. Government could, for example, employ chat bots to improve service delivery in areas where processes could be easily automated and do not have a high risk of causing harm.

Some speakers highlighted how AI could be harnessed to support government decision making. We heard from Michael Karlin, Senior Advisor at the Treasury Board of Canada Secretariat, that the federal government is exploring how to potentially use AI to improve recordkeeping, identifying documents with business value from those that are merely transactional. Jimoh Ovbiagele, Co-founder and Chief Technology Officer at ROSS Intelligence, noted that policymakers could use NLP for social media sentiment analysis in order to better understand citizen reactions to policy in near real time.

Michael Karlin stressed that government should only apply AI where it makes sense - that is, where it provides benefit and can be used responsibly. Lex Gill suggested this necessitates that government reflect upon current systems, identify their shortcomings, and determine where AI could be used to improve upon these gaps. When an intelligent system is used to make decisions that directly impact the lives of citizens, speakers cautioned that government has the responsibility to ensure the technology performs accurately and fairly, informed by ethical guidelines and values, and without perpetuating existing systemic biases and deficiencies.

Government as a collaborator

Nearly all speakers highlighted the need for government to take a collaborative approach in regards to regulating AI, both internally within and across government ministries, and externally, with academia and industry. Many speakers urged that no one central ministry can, or should, be tasked with addressing AI; government responses to AI must be broad-based and coordinated. As Michael Karlin said, the “nature of confederation makes

things complicated but it is actually our biggest strength, because we have lots of different innovation going on at different levels of government, and the best thing we can do is convene everyone and learn from each other.” Alan Veerman, Chief Operating Officer at the Vector Institute for Artificial Intelligence, noted that public policy is naturally cross-cutting, which can potentially slow down initiatives. He described responding to and incorporating AI as the biggest change management exercise ever undertaken by major companies and institutions.

The need for external collaboration with industry and academia was a common theme throughout the day. Speakers from CIFAR, the Ontario Cabinet Office, Vector Institute, and the Treasury Board of Canada all spoke of the value of multi-sectoral collaboration, the opportunity for government to learn from industry and academia, as well as the need for multi-disciplinary discussions. Abhishek Gupta highlighted the need for government to have inclusive discussions in order to define the legal and moral boundaries of AI systems.



*Pictured: Dr. Avi Goldfarb,
“The Simple Economics of AI”*



*Pictured: Dr. Kathryn Hume,
“Historical Context – Why Now?”*

FACILITATED DISCUSSIONS

Themes from the facilitated discussion are reflected below, along with questions these discussions generated for further research and analysis. As AI was a new field for many participants, the discussions reflect many of the issues presented by speakers and are not exhaustive of opportunities and challenges, or necessarily fully up-to-date on the legal and regulatory environment of AI or technological advances. Many opportunity areas had a corresponding challenge, demonstrating the dynamic potential of AI to do both harm and good. For example, participants discussed the potential to automate judicial decisions to perpetuate existing biases in the legal system, reflecting both policing biases and over-surveillance in particular communities and historic bias in sentencing. But they also discussed the potential for automated justice to recognize and adjust to mitigate this, eventually becoming less biased than the average human judge.

VALUES

Participants discussed how technology can reflect the values of the people who create it and expressed concerns about the diversity and inclusion currently represented within the tech sector and the power dynamics between designers and users. Some participants noted that AI systems have the potential to control or limit the types of opportunities that people have access to, and that AI technology can, unconsciously or consciously, be designed in ways that have a high potential to perpetuate discrimination and oppression. This theme was particularly strong among the groups exploring Automated Hiring, which exemplified the potential for algorithms based on data collected on current employees to encourage hiring of people with similar backgrounds. They also discussed the use of AI for end goals that may not reflect Canadian values and norms, such as mass persuasion (or what Abhishek Gupta referred to as “attention hacking”), surveillance, and data collection.

Key questions included:

- + Whose values are represented or served by AI applications?
- + How are these values being built into and reinforced by AI systems?
- + How should these systems be assessed and evaluated and who should set the benchmarks?

BIAS AND DISCRIMINATION

Across all discussions, participants recognized the existence of bias as both a feature of our current social and governance systems, and a design flaw within AI technologies, and expressed concern that the impacts of this bias would be unevenly felt by historically marginalized groups. Lex Gill noted that training data that is predominantly images of certain groups and ethnicities will inevitably reinforce system forms of discrimination and a false path to dependency. Participants also discussed the potential for AI to mitigate human bias, enable fairer, more equitable public services, and improve decision-making capabilities within government as well as private companies (e.g., using algorithms to assess and address systemic bias in hiring).

Key questions included:

- + What biases currently exist in our social and governance systems? Who is most vulnerable in an AI-enabled society?
- + How can we prevent AI systems from inheriting and building biases? What kind of governance or policies should be in place?
- + Could AI be less biased than humans? How can AI be used to mitigate bias in human decision-making?

DATA OWNERSHIP AND CONTROL

Participants expressed significant concerns about the ownership and security of personal data, and the unequal distribution of benefits from data that is already being collected. They discussed the potential for surveillance by corporations operating between or outside of national boundaries and regulation and suggested that governments explore new frameworks for privacy, consent, and data ownership.

Key questions included:

- + Who owns metadata (data that contains information about other data, such as the time, date, and location of a Facebook status)?
- + In the context of data collection, is individual consent (for data collection, sharing, and use) the right framework?
- + How can consumers be protected from malicious private interests (e.g., social engineering, political propaganda, and attention hacking)?

TRANSPARENCY AND OPEN SOURCE ALGORITHMS

In his lightning talk on accountability, Michael Karlin stressed the importance of transparency in government AI in order to explain to Canadians why and how a decision was made, and the role that transparency plays in enabling public discourse and democratic challenges to government decisions. In AI, transparency usually refers to the ability to see and audit algorithms, data sources, and automated decisions. Prompted by this talk, participants debated whether government use of AI should be transparent by default or whether there are some kinds of decisions that don't require transparency. Participants expressed interest in an assessment of, or criteria for, government transparency. They asked: does the AI system that

determines traffic light signals require the same transparency and oversight as a system that determines student loan eligibility? Some participants suggested that government has an opportunity to lead by example by using open source algorithms, setting standards on transparency and creating code that can be repurposed. Open source AI was also identified as an opportunity for citizen empowerment, offering the means for civil society to understand, assess, and challenge government decisions. Organizational silos within government were also identified as a potential challenge to the roll-out of AI for service delivery and operations, including difficulties of collaborating and sharing information and data across ministries.

Key questions included:

- + What should the criteria be for making automated decisions transparent?
- + If an individual believes the automated decision was made in error, what recourse is available? Should this process be different for government use of AI versus private sector use of AI?
- + How could governments do a better job of making data available when using AI to tackle complex policy issues?

PRIVACY

Participant discussions of privacy were rooted in the idea of privacy as a public good, introduced by Lex Gill's lightning talk. Participants explored how government could apply differential privacy – the ability to collect, aggregate, and spot patterns in data without compromising individual privacy – to protect personal information, and discussed public expectations around privacy and the challenge of defining its legal and social boundaries. They voiced concerns about the need to protect citizen

safety and privacy against AI-enabled hacking, while recognizing the potential for AI to be used to enhance cybersecurity. While privacy was largely considered to be a benefit, some participants felt that increased privacy could present a challenge to accurate representation in data sets, potentially leading to individuals experiencing incorrect or biased outcomes. Government was seen as having a higher prerogative to ensure privacy of personal sensitive information than their private sector counterparts. However, participants noted the challenge for government to define the legal and social boundaries of “privacy”.

Key questions included:

- + Should privacy be a public good?
- + Could citizens benefit from government use of differential privacy?
- + What are the differentiated contexts of privacy? What are Canadian norms and expectations around privacy?

“The consent model is individualistic, transactional, contractual, and it’s really only adequate if we want to think about data as a commodity that we give and take rather than as embedded with rights that we have. Instead I think it might be helpful, as some scholars have proposed, for policymakers to think about privacy as a public good...like clean drinking water or a functional healthcare system.”

- Lex Gill, Research Fellow at Citizen Lab

GOVERNMENT REGULATION AND CONSUMER PROTECTION

Workshop participants identified the valuable role of government intervention and oversight in responding to the threats and opportunities of AI and ensuring public interest and protection of values. Some participants highlighted the need for explicit regulation of the industry, including updating or developing laws regarding data collection, privacy, transparency, accountability, and what decisions can be automated. Others noted the need for cross-sectoral collaboration in developing these policies, to ensure that they reflect current technologies and applications, and cutting-edge research.

Key questions included:

- + How will current legal/governance frameworks be applied to AI? Where are the gaps in regulation and policy?
- + How can government work proactively and collaborate to protect public interest? What is the best way to ensure industry compliance?
- + What amount of technical knowledge or detail is needed to effectively think through how and where AI could be applied in policy making?
- + Is there a role for a federal AI agency or other national body to investigate AI’s impact on society?

NEXT STEPS

By bringing together a range of participants from government, industry, and academia, this event provided a space for dynamic, multi-disciplinary discussions on the cross-cutting challenges and public policy implications of AI. The issues we explored were technically, legally, and socially complex and a number of participants indicated that the event would have benefited from a two-day format and more time for participant engagement, or future deep-dives into key issues such as regulation, data sovereignty, privacy, and AI talent development. BII+E looks forward to supporting future discussions, collaboration, and research on AI and public policy.



Pictured: “Decoding AI: Impacts for Policymakers”, Facilitation Exercise

APPENDIX 1: AGENDA

Morning: Level-Setting

Time	Event
8:00 am	Registration Opens, Networking + Light Breakfast
8:45 am	Opening Remarks from the Government of Ontario + the Brookfield Institute <i>Sean Mullin, Brookfield Institute for Innovation + Entrepreneurship</i> <i>Jay Porter, Policy Innovation Hub, Government of Ontario</i>
9:00 am	Historical Context: Why Now? <i>Dr. Kathryn Hume, integrate.ai</i>
9:20 am	AI 101: Understanding the Technology – Part I <i>Dr. Graham Taylor, NextAI + University of Guelph</i>
9:40 am	AI 101: Understanding the Technology – Part II + Predictive Analytics – <i>Brian Purcell, IBM</i> + Natural Language Processing – <i>Jimoh Ovbiagele, ROSS Intelligence</i> + Computer Vision – <i>Dr. Kosta Derpanis, Ryerson University</i>
10:20 am	Break
10:50 am	So What? Implications of AI for Policymakers <i>Dr. Elissa Strome, CIFAR</i> <i>Michael Karlin, Treasury Board of Canada Secretariat</i> <i>Alan Veerman, Vector Institute</i>
11:45 am	Lunch + Networking
12:30 am	Special Topic: “Digital Deceit” <i>Dr. Taylor Owen, University of British Columbia</i> <i>Dr. Ben Scott, New America</i>

Afternoon: Drilling Down

Time	Event
1:00 pm	<p>Cross-Cutting Implications of AI</p> <ul style="list-style-type: none">+ Ethics – <i>Abhishek Gupta, District3</i>+ Bias + Privacy – <i>Lex Gill, Citizen Lab</i>+ Explainability – <i>Michael Karlin, Treasury Board of Canada Secretariat</i>+ Safety – <i>Alexandre Piché, Montreal Institute for Learning Algorithms</i>+ Accountability – <i>Carole Piovesan, McCarthy Tetrault</i>
2:30 pm	Break
2:45 pm	<p>The Simple Economics of AI <i>Dr. Avi Goldfarb, University of Toronto</i></p>
3:00 pm	Decoding AI: Impact for Policymakers
4:15 pm	Wrap Up + Reflections
4:30 pm	End of Day

APPENDIX 2: CASE STUDY DISCUSSION PROMPTS

Case Study :



If you picked your own topic tell us why?

<p>What inspires you about this topic?</p>	<p>Potential application in your policy area? Consider programs users, data, techniques and resources.</p>
<p>How will AI help?</p>	
<p>What other policy domains/jurisdictions will be impacted by this topic?</p>	<p>Potential implications for Canadians - Are there any vulnerable groups to be considered?</p>

Considerations? Opportunities and Challenges		List top 3 areas for further exploration & possible public policy
Bias		<div style="display: flex; align-items: center; margin-bottom: 10px;"> 1 <hr style="border: 0; border-top: 1px solid #ccc; width: 100%;"/> </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> 2 <hr style="border: 0; border-top: 1px solid #ccc; width: 100%;"/> </div> <div style="display: flex; align-items: center;"> 3 <hr style="border: 0; border-top: 1px solid #ccc; width: 100%;"/> </div>
Explainability		
Safety		
Privacy		
Ethics		
Accountability		

APPENDIX 3: CASE STUDY CARDS

Automated Hiring

Automated assessments of job applications are becoming increasingly popular. In-person interviews are recorded on camera, using computer vision, algorithms, and other technologies. The software assesses barely perceptible changes in posture, facial expression and vocal tone, turns the data into a score, and compares these to data collected from existing top-performing employees. Though in some jurisdictions, data protection regulation requires companies to disclose whenever a decision that significantly affects an individual has been automated, minimal human involvement (such as approving a list of automatically-ranked CVs) could exempt companies.

Smart Homes

A housing developer intends to include smart home technology driven by artificial intelligence in a large condo complex under development. This technology includes smart security cameras, built-in voice-activated home systems (e.g., Google Home, Alexa, etc.), and doorbells with voice and facial recognition. Residents will be unable to opt-out of the installation or data collection by this private company and their data could potentially be sold or shared with other organizations, including law enforcement and the government.

Autonomous Vehicles

You are checking your emails while your car operates on autopilot. Outside, the wind picks up and it begins to rain heavily. Lightning strikes a tree ahead of your car. It falls onto the road, directly in front of you. Before you have time to assess the situation and take control, the car, trying to avoid the tree, swerves and hits a cyclist to your right, knocking them off their bike. They are taken to the hospital with life threatening injuries.

Chatbots

A telehealth application that provides for-fee real-time digital healthcare rolls out across Canada. After filling out a registration form with personal information (e.g. name, age, sex, weight, height, smoker/non-smoker), the user is introduced to a chatbot which asks a series of questions to filter out possible causes or illnesses. Once the user has answered each question sufficiently, the chatbot provides them with a possible diagnosis. The user's answers and diagnosis are stored in a cloud belonging to the telehealth company, and saved for future reference if that individual were to use the app again. For an additional fee, users can chat remotely with a live doctor.

Predictive Life Insurance

A life insurance company purchases health data from fitness applications (e.g. Fitbit, MiCoach, Strava). It implements a predictive algorithm to determine insurance premiums based on the data collected. Individuals are scored based on the amount of healthy activity they engage in (e.g., number of steps taken per day, heart rate level, sleep duration). Individuals with high scores are awarded lower premiums than those who engage in less activity and get less sleep.

Road Maintenance

The Ministry of Transportation implements a new system for scheduling road works across the GTA. The system collects data on road conditions in the area, using a combination of crowd-sourced reports from citizens via a mobile app and 311 calls, ministry records, and on-the-ground surveys. An algorithm is then used to analyze this data, prioritize work based on importance, and schedule road works at times of low volume or traffic.

Employee Tracking

A manufacturing company requires employees to wear smart bands, internet-connected bracelets equipped with GPS and activity tracking that interacts with sensors throughout the factory. Throughout the work day, the smart band monitors every employees' location within the factory and the amount of time they spend on each task, including washroom breaks. The data generated by this wearable device is recorded and analyzed by a machine learning algorithm, enabling management to see patterns in employee behaviour and modify business processes to create maximum human efficiency. It also enables employers to rank employees in order of productivity.

Adaptive Education

The Toronto District School Board decides to implement an AI assistant in K-12 education. The AI assistant tracks each student in real-time, identifying their strengths, weaknesses, and learning styles. When a student graduates from one grade to the next, the data collected about them is carried over and continues to inform their learning experience. The AI assistant augments the traditional teaching environment by providing each student with a tailored set of modular lessons. It is capable of identifying when a student becomes bored or distracted and can adapt to each individual's changing levels of knowledge and needs. Importantly, it alerts teachers and parents of students who need extra help or are falling behind. Overtime, the system automatically begins to classify students into categories, such as (1) Likely/Unlikely to Graduate, (2) Likely/Unlikely to Attend University, (3) Likely/Unlikely to Break the Law.

Prescriptive Legal Analytics

Judge Insights, an online platform, collects and analyzes historical behavioural information about Canadian judges, and offers lawyers predictions on how a particular judge would act in a particular court case. Lawyers fill out an online template with information about their client and the case, and the judge's past decisions, and the algorithm provides likely outcomes and recommended approaches. In doing so, lawyers can gain insight into the best course of action for presenting evidence in court, convincing the judge, and influencing the jury.

Detecting At-Risk Youth

The Ministry of Children and Youth Services adopts a profiling system to identify at-risk youth in the foster care system. The system employs a machine-learning algorithm that produces a risk score based on a number of variables including the neighbourhood they live in, the school they attend, their families and peers, and their interactions with social services and police. The more risk factors associated with a child, the higher their risk score, and the more likely it is that the system will recommend intervention.

APPENDIX 5: PARTICIPANT EVALUATIONS

Sixty-eight participants attended, including a mix of academics, industry (both large companies and AI startups) and public servants representing 16 ministries across the Ontario and federal governments. Twenty-five participants filled out post-event surveys (36 percent of attendees). Participants were asked to rate the event on a scale of 0 (low) to 5 (high):

Would attend a related event: 4.88 / 5

Organization: 4.76 / 5

Good Use of Time: 4.68 / 5

Morning presentations were engaging: 4.42 / 5

Afternoon sessions were well facilitated: 4.64 / 5

Clear Workshop Goals: 4.36 / 5

Has today's event changed the way you think about AI in the context of policy making? 92 percent of respondents said "Yes"

What is the most surprising thing that you learned today?

- + The diversity and depth of policy issues and the immediacy of challenges posed by AI
- + That the provincial civil service cares deeply about the opportunities presented by AI
- + Conceiving of AI as a predictive process and a tool able to reduce the cost of prediction

What was the most valuable component of today's event?

- + In-person interactions: talking to speakers about applications to participants' work; networking and in-person connections to people working on AI
- + Access to a diversity of perspectives on regulating and applying AI
- + In-depth discussions during the workshops following the lightning talks; getting to apply the concepts
- + Hearing from CIFAR, Vector Institute, and Dr. Ben Scott

What was the least valuable component of today's event?

- + Some participants reported that the introductory morning sessions were the least valuable event component. Some responded that they were too long, and covered material already presented in the background package.
- + Other participants reported that the sessions were too overwhelming and contained too much information.

Is there a topic that you wanted us to cover today that we may have missed?

Participants responded that they wanted more deep dives and more time to explore the issues presented. They proposed exploring:

- + How AI is currently being deployed and a critical evaluation of the status quo
- + Accountability and the need for tech literacy
- + The impact on labour
- + Transnational regulatory coordination
- + Public sector investment in AI
- + How the public sector has grappled with previous emerging technologies
- + Potential policy responses

What suggestions do you have for future events related to this topic?

- + A number of participants suggested that future events could be two days rather than one, allowing time for in-depth discussion of issues such as ethics, accountability, regulation, the use of AI for harm, data sovereignty, and AI talent development.
- + Some also proposed more focused events on specific topics (e.g., transparency, data ownership) and more time for interactive workshops and participant involvement.