

AI Futures Policy Lab: Ottawa

January 2019



AUTHOR



SARAH VILLENEUVE
Policy Assistant

Sarah conducts research under the AI + Society and Innovation + Inclusive Economic Growth work streams. She is a member of 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 the University of London.

ACKNOWLEDGEMENTS

This engagement is part of a series of workshops hosted by the Brookfield Institute for Innovation + Entrepreneurship (BII+E) and the Canadian Institute for Advanced Research (CIFAR).

CIFAR brings together outstanding researchers from across disciplines and borders to address important challenges facing the world. We support leading edge research with the potential for global impact.

Established in 1982, CIFAR is a Canadian-based, international research institute with nearly 400 fellows, scholars and advisors from 18 countries. In 2017, CIFAR was chosen by the Canadian government to lead the \$125 million Pan-Canadian Artificial Intelligence Strategy.

For more information, visit cifar.ca.



MaRS Centre, West Tower
661 University Ave.,
Suite 505, Toronto, ON
M5G 1M1

CONTRIBUTORS

Brent Barron, Director, Public Policy, CIFAR
Karen Birkemoe, Graphic Designer
Gaga Boskovic, Public Policy Associate, CIFAR
Meghan Hellstern, Senior Projects Officer, BII+E
Heather Russek, Director, Policy Innovation Platform, BII+E
Jessica Thomson, Digital Content and Marketing Coordinator, BII+E
Erin Warner, Marketing and Communications Specialist, BII+

The Brookfield Institute for Innovation + Entrepreneurship (BII+E) is an 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.



20 Dundas St. West, Suite 921
Toronto, ON
M5G 2C2

In January 2018, CIFAR and the Brookfield Institute for Innovation + Entrepreneurship (BII+E) formed a partnership to design and host five AI Futures Policy Labs aimed at generating greater awareness of the long-term implications of AI and exploring the future of AI policy in Canada. Between June and October 2018, CIFAR and BII+E held labs in Toronto, Edmonton, and Vancouver, with the participation of 63 emerging policy leaders.

On November 22, 2018, CIFAR and BII+E hosted the fourth AI Futures Policy Lab in Ottawa, Ontario. This event brought together 28 emerging policy leaders with the aim of:

- + Building capacity of future public service leaders to understand the policy implications of AI and respond appropriately
- + Providing policymakers with a direct line of sight into the AI sector: the myths and hype, the evolving state of technological advances, and potential applications
- + Contributing to early government responses to emerging AI technologies

To achieve these aims, this lab was designed to raise awareness of the opportunities and challenges associated with current AI capabilities and applications, encourage critical thinking around potential future scenarios, and facilitate the development of policy recommendations. Feedback from the previous three AI Futures Policy Labs were used to re-design certain aspects of the lab's agenda and content. Participants were presented with a case study featuring a current AI application associated with a specific policy domain (i.e. housing, justice, education, immigration, and hiring). Facilitators guided each respective group through their case study and accompanying prompts. During the final session of the day, groups presented policy recommendations related to the opportunities and challenges associated with their case study. The agenda developed for the day is provided in appendix A.

CASE STUDY POLICY DOMAINS

Prior to the lab, five sets of case studies were developed. Each set was associated with a specific domain: housing, justice, education, immigration, or hiring. Participants were organized into groups of 4-5 and assigned to a domain.

HOUSING

AI is impacting the housing sector in multiple ways, from smart-home devices like [Nest](#) to intelligent tools that help to curb energy use, and services that even act as the middle-man between landlords and tenants. Advancements in this domain afford residents with potential benefits, but also create challenges regarding privacy and safety in a domestic environment. Within this domain, participants examined [Naborly](#) (appendix B), a tenant screening application that generates risk scores to help landlords make smarter letting decisions.

JUSTICE

The legal sector is being impacted by recent developments in AI and machine learning capabilities that have enabled applications to automate legal research, due diligence processes, contract review and management practices, and help to predict legal outcomes. Participants within this domain were given the chance to explore the policy impacts of [ROSS intelligence](#) (appendix C), an artificially intelligent legal research tool that applies natural language processing to increase lawyer's ability to identify relevant information.

EDUCATION

There is vast potential for AI to transform education in ways that make learning more accessible, provide personalized curriculum, and support educators in delivering content. Participants in this group analyzed [Nestor](#) (appendix D), an artificial intelligence class assistant that uses machine learning algorithms and advanced facial recognition to analyze the attention of students listening to online lectures.

IMMIGRATION

Participants within this group examined an AI application that is automating Canada's immigration process¹ by sorting applications into two streams: simple or complex (appendix E). This effort has been undertaken by Immigration, Refugees, and Citizenship Canada (IRCC) in an attempt to ease the backlog of immigration applications that immigration officers are faced with. Applications deemed as simple do not need to undergo review, and are processed at a faster rate than those that are identified as complex. Complex cases must be reviewed by a human and face longer processing times.

HEALTH

A number of large companies within Canada are integrating AI screening applications into their hiring processes. Participants within this group examined [Ideal](#) (appendix F), a talent intelligence application that centralizes data gathered from applicants' resumes, chatbot conversations, and online assessments to screen and analyze candidates in real time. Ideal then identifies and provides the employer with a shortlist of strong candidates.

POLICY LAB ACTIVITIES

1. THE 'AI' THING FROM THE FUTURE

The lab began with an ice-breaker game, The 'AI' Thing From The Future.² The purpose of this activity was to encourage participants to be creative and think beyond the current reality. The format of this game was changed from the previous labs because a new edition of the original game with updated card prompts was released. CIFAR and BII+E also added an AI card to the prompts to encourage participants to think about specific AI capabilities and applications. Each table of participants played two rounds of The 'AI' Thing from the Future with the help of a facilitator. Each group was then given four cards, each containing a unique prompt related to the type of city this future is situated in, object of focus, a theme, and an AI capability or application. Participants were provided with a template to record their ideas.

2. AI 101

Dr. Joel Martin, Director of Research & Development at the National Research Council of Canada, kicked off the speaking sessions by providing participants with an introduction to AI capabilities and applications. His presentation offered an overview of machine learning (including deep learning and reinforcement learning), computer vision, and speech recognition, and relevant applications to energy, language translation, wearable technology, and gaming. Dr. Martin suggested that the next splash in AI may be based on increased computation power or the combination of currently distinct capabilities, such as natural language processing and computer vision to, for example, learn a language.

¹ Adapted from "Bots at the Gate: A Human Rights Analysis of Automated Decision-Making in Canada's Immigration and Refugee System" © 2018 International Human Rights Program (Faculty of Law, University of Toronto) and the Citizen Lab (Munk School of Global Affairs and Public Policy, University of Toronto). Licensed under the Creative Commons BY-SA 4.0 (Attribution-ShareAlike Licence)

² Adapted from Stuart Candy and Jeff Watson (Situation Lab)

3. AI AND PUBLIC POLICY IN CANADA

Michael Karlin, Senior Advisor at the Treasury Board of Canada Secretariat, spoke about the opportunities AI offers to the Canadian public service, and the current efforts underway to effectively harness this technology. He also highlighted the challenges AI poses to government agencies looking to utilize this technology. For example, he addressed issues concerned with procedural fairness. This included the Treasury Board of Canada Secretariat's [Algorithmic Impact Assessment](#) and [Directive on Autonomous Decision-Making](#), which are being developed to ensure the responsible adoption of AI within the Government of Canada. Karlin also spoke about the CIO Strategy Council's initiative to develop standards for ethical AI and data privacy, and flagged that low support within government and lack of citizen awareness hamper efforts to integrate this technology.

4. AI AND POLICY

Karlin's talk was complemented by a short presentation on the broader picture of the AI policy landscape in Canada and abroad, presented by Brent Barron, Director of Public Policy at CIFAR. This presentation provided an overview of the [Pan-Canadian AI Strategy](#), federal supercluster funding (e.g. [SCALE.AI](#)), and the [Montréal Declaration on Responsible AI](#), as well as provincial initiatives in British Columbia and Alberta. This overview also highlighted international examples, such as the General Data Protection Regulation (GDPR) in the European Union, the introduction of the AI in Government Act in the United States, the British government's investment in skills and commitment to be a leader in ethical data use, and China's industrial AI strategy.

4. ANALYZING CURRENT AI APPLICATIONS

Within their pre-assigned groups, participants were presented with an example of a current AI application (Naborly, Nestor, ROSS Intelligence, Ideal, or AI for Immigration). Each group was given time to read the case study and discuss any preliminary questions within their groups before turning to the canvas. Once the group was comfortable in their understanding of the case study, they turned their focus to the first canvas (appendix G). This canvas prompted participants to think about the types of individuals or groups that are impacted within this case study (positively and/or negatively); the potential impacts of this technology at the local, national, and global levels; and existing policies and programs affected by the technology.

Facilitators encouraged participants to actively contribute by writing their thoughts on sticky notes and placing them on the canvas, first individually and then interactively as a group.

5. IMAGINING AI IN 2028

Facilitators then led their groups into an open discussion about how their case study may look in 2028. This discussion has proven useful in enabling participants the freedom to imagine future scenarios that incorporate the same themes of their current AI case study. This included imagining how each application may develop and impact individuals, communities, and policies, as well as social, cultural, political, and economic processes within the next 10 years.

6. TAKING ACTION TODAY

Following this discussion, facilitators presented their groups with the second canvas (appendix H), which prompted participants to reflect on the discussions from the previous two sessions. While they reflected, participants were asked to think of which individuals and/or groups would experience the most notable effects (both positive and/or negative), as well as highlight the most significant socio-political effects (both positive and negative) pertaining to the current AI case study. Each group was provided with a template (appendix I) that asked them to write a description of the case study they examined, the associated opportunities and challenges, and their top 3 policy recommendations.

POLICY RECOMMENDATIONS³

DOMAIN: HOUSING

Participants in this group identified contradictory benefits and challenges associated with Naborly. The first of these relates to bias. While participants noted that the application could reduce pre-existing biases related to things like ethnicity or socioeconomic status, they also recognized that new sets of biases could emerge as the system learns from data. Secondly, while the application streamlines the submission of applications from prospective tenants, it may also contribute to accelerating gentrification by giving applicants with stronger employment, credit, and rental histories a higher ranking than those who are lacking in these areas. Thirdly, participants recognized that although Naborly offers a more efficient way to manage and access records, it also raises concerns over data security.

Recognizing that there are jurisdictional challenges to addressing the risks posed by this application, the group provided the following recommendations:

- + Hold public consultations about AI and data to raise awareness and address concerns about privacy and security.
- + Establish a Department of Digital Policy and an Office of the Algorithmic Intelligence Auditor to manage government response in an agile manner as technology changes.

DOMAIN: JUSTICE

As a tool that can assist lawyers with identifying relevant information and tracking changes in case law, participants highlighted the potential for ROSS Intelligence to improve efficiencies in the justice system and reduce the cost of legal services, resulting in positive implications for inequality. The equality-enhancing effects of AI-driven legal research are dependent on ensuring that the technology is broadly available, and not just employed at large firms. The group also acknowledged the risk of lawyers becoming over-reliant on systems like ROSS, as well as the potential for bias to creep into the system. These risks could result in a loss of trust, not just in the application, but in the justice system as a whole. Finally, the group explored the possibility that the decreasing cost and increasing efficacy of automated legal research could lead to a strengthening of the role of precedent, since there will be a greater volume of relevant case law identified.

Policy recommendations:

- + Community-access strategy that includes procuring autonomous research tools in legal aid and similar organizations, the creation of a community feedback liaison that can connect community responses to technology developers, and open data.
- + Legal profession reforms, such as curriculum updates, professional standards, code of conduct.
- + Grassroots Innovation Challenges: Hackathons to identify and mitigate bias in justice applications.

³ Disclaimer: The following policy recommendations were developed by participants through an exercise designed to help emerging policymakers explore existing policy levers in relation to specific case studies. These do not represent the views of CIFAR and BII+E.

DOMAIN: EDUCATION

As an application that aims to enhance student learning experiences and outcomes, participants valued Nestor's potential to facilitate more personalized and inclusive learning. They also highlighted other potential applications of the underlying technology behind Nestor, namely using its ability to detect attention and send reminder messages in contexts like driving. They also recognized the larger beneficial impact this kind of application could have on economic development, including the ability to support the growth of Canadian AI firms.

Policy recommendations included:

- + An AI framework and principles to guide use in education.
- + An innovation fund to support adoption by universities and colleges to implement Nestor.
- + Educational non-profit data trust to protect the personal data collected by educational AI applications.

DOMAIN: IMMIGRATION

This group recognized both the potential benefits and risks of automating the immigration screening process. One benefit participants highlighted was the tool's ability to serve as a mechanism for reducing the current immigration processing backlog and the potential to scale up this technology for other internal purposes. This group also identified a number of challenges introduced by this tool. These included the potential for bias in the design of the system, leading to a lack of procedural fairness and a potential increase of discriminatory outcomes. The group also discussed how these risks could lead to larger-scale economic and geopolitical concerns, such as potential market and labour shifts resulting from workers migrating across borders at faster rates. The group was most concerned about inaccurate use of this system ultimately contributing to unpredictable large-scale societal changes, both on national and global scales.

Policy recommendations:

- + Establish an independent oversight body with a phase-gate implementation process.
- + Ensure system is designed and used in alignment with human rights legislation.
- + Utilize the Treasury Board of Canada Secretariat's Directive on Automated Decision Making to provide guidance on the deployment of this technology.

DOMAIN: HIRING

Participants in this group described the many benefits Ideal has to offer, in particular for large companies that have a high turnover rate. These include the ability to reduce personal hiring manager bias by having a program screen and rank candidates. Ideal also reduces the time it takes for companies to identify and create a shortlist of strong candidates, which in turn reduces cost. Participants also noted the application's ability to identify employees who are a better fit for the job, and therefore more productive. However, they discussed risks such as bias in the automated screening process, exacerbating income inequality, and potential for abuse or misuse by employers.

Policy recommendations:

- + Create an arms length research institute to study and disseminate best practices on hiring and AI.
- + Develop an AI explainability framework to make decisions transparent
- + Make an open-source version of this system to improve transparency.

GENERAL REMARKS

The participants at the Ottawa lab emphasized the benefits of having the opportunity to learn about and examine real-life applications. This allowed them to gain a better understanding of the underlying technology and helped to ground the day's discussions. Participants also valued the policy recommendation portion of the day, as it enabled them to synthesize their discussions and think through relevant policy responses. As with previous labs, Ottawa participants enjoyed the opportunity to network. However, participants noted that they would have liked more interaction between groups throughout the day to hear about what everyone was working on, and potentially contribute to each other's discussions. Participant feedback also indicated a desire for more interdisciplinary discussions on current capabilities and applications of AI, as well as stronger government responses.

NEXT STEPS

CIFAR and BII+E will be reviewing participant feedback from all previous labs and iterating on the design of our final AI Futures Policy Lab taking place in Montreal, Quebec in early 2019. If you would like to participate in our upcoming lab in Montreal, or other related events, please contact [Gaga Boskovic](#). A final report summarizing all five in this series of AI Futures Policy Labs will also be published in early 2019.

APPENDICES

APPENDIX A: AGENDA

Time	Activity
8:30am	Light Breakfast + Networking
9:00am	Opening Remarks <ul style="list-style-type: none"> Brent Barron, Director of Public Policy, CIFAR Heather Russek, Director, Policy Innovation Platform, The Brookfield Institute for Innovation + Entrepreneurship
9:15am	Activity: Thing From the Future
9:35am	AI 101: Dr. Joel Martin, Director R&D, National Research Council Canada
10:20am	Break
10:35am	AI and Public Policy: Michael Karlin, Senior Advisor, Treasury Board of Canada Secretariat
11:15	AI Policy: Brent Barron, Director of Public Policy, CIFAR
11:30pm	Lunch
12:30pm	Activity: Analyzing Current AI Applications In small groups, participants will dive deeper into a current application of AI, analyzing its social, economic, and political impacts. Groups will also be asked to forecast what this technology might look like in a year, and what new implications this may have.
1:30pm	Discussion: Examining AI in 2028 In small groups, participants will explore the social, political, economic, and ethical dimensions of future AI scenarios.
2:00pm	Break
2:15pm	Activity: Taking Action Today Reflecting on the previous sessions, participants will brainstorm relevant forms of government interventions that can be used to support the ethical development and beneficial use of AI. In small groups, attendees will collaboratively draft a short policy recommendation based on the case studies that have examined throughout the day.
3:15pm	Activity: Presentations & Closing Each group will have the opportunity to present their policy recommendation to the larger room and reflect on the day.
4:00pm	Social & Networking (Off-Site)

APPENDIX B: NABORLY

Founded in 2015, Naborly is a tenant screening application that generates risk scores, enabling landlords to make smarter letting decisions.

Naborly serves as a free online application for property rentals. Landlords send prospective tenants a link to the online application to fill in their rental history, employment, and financial information. Naborly then analyzes and produces an applicant risk score based on the applicant's income, identity and employment, credit ratings, criminal records, and rental history. Naborly's Applied Artificial Intelligence system, SHERLY, an inductive, deductive, and reductive reasoning system, continuously learns from thousands of rental applicants and their tenancy outcomes, allowing it to better identify patterns of risk.

Through this process, Naborly removes traditional factors of discrimination stemming from landlord biases relating to tenant finance, social class, or race. Additionally, Naborly adjusts its scoring for each unique applicant, taking current rental property characteristics and the market prices into consideration. Results are delivered to the landlord within minutes of the application being submitted.

After an application is submitted, Naborly automatically creates a personal private profile, which stores information for future applications. This helps tenants build a verified rental history without the need for printing, scanning, and faxing documents. All information collected by Naborly on prospective and current tenants, landlords, as well as API Partners is protected by a state-of-the-art data security infrastructure. This ensures that the data held by Naborly remains accessible only to authenticated users and recipients with expressed permission from the user. Tenants can then use Naborly to apply to landlords that do not yet use the system.

Naborly democratizes rental record keeping through the use of its global open_DOOR database system, which allows tenants, landlords and property managers to share feedback, evictions, judgements, and verified disputes. This provides both prospective applicants and tenants with an added layer of transparency before entering into a rental contract. While Naborly is fully compliant with Privacy and Fair Housing laws across the US and Canada, and its algorithms are regularly audited to ensure it continues to meet the requirements for compliance, this does not mean its system is verified beyond its compliance to these laws.

APPENDIX C: ROSS INTELLIGENCE

ROSS is an artificially intelligent legal research tool that applies cutting-edge natural language processing (NLP) to increase lawyer's ability to sort through and find information relevant to their cases. Lawyers need to do substantial legal research to prepare for a case, normally taking days, weeks, or even months to source out information - but ROSS can now automate this process. Using a combination of advanced keyword search and machine learning, ROSS enables lawyers to identify relevant information faster and more efficiently, and even uncover information that could have been missed by sifting through over a billion text documents per second.

ROSS's advanced NLP technology has been trained to understand legal jargon and encompasses all American case law. Lawyers can enter queries such as, "When is secondary liability with respect to copyright infringement established?" and receive an overview of relevant key points drawn from a database of published and unpublished case law, substantive law, procedural law, and legal analysis.

ROSS is also able to track relevant developments in the law related to a specific legal issue and notify lawyers of relevant legal updates. Additionally, lawyers are able to upload a range of legal documents, such as memos, motions, or briefs, for ROSS to analyze and flag cases cited in the document that have received negative treatments in court.

Built on IBM Watson's cognitive computing platform, ROSS learns from past interactions and improves its accuracy the more its system is used. ROSS is currently used by law firms such as Baker Hostetler and Latham & Watkins LLP.

APPENDIX D: NESTOR

Nestor, developed by LCA Learning, is an artificial intelligence class assistant that uses machine learning algorithms and advanced facial recognition to analyze student attention while listening to online lectures. The software is currently being used for two online courses offered through the ESG business school in Paris, France.

Nestor aims to enhance the performance of both the student and the teacher. Using students' webcams, Nestor's facial recognition software tracks 20 key landmarks on the students face - including the eyes, brows, mouth, and jaw - and can even detect when a student has pulled out their phone. Facial expressions are measured using three variables. The first is engagement, which measures facial muscle activation that detects expressiveness and responsiveness. The second is valence, which measures the positive and negative facial expressions. The third is attention, which measures focus according to head orientation.

Once the system detects the student has lost focus, it can send a message alerting them to pay attention. Nestor can also predict when a student may start to drive away again, sending them a signal to stay focused before attention is lost. Nestor also quizzes students on content that was covered while they appeared to be distracted. Student performance and attention analysis, particularly when focus decreases, is then relayed to the teacher who can adjust future lessons appropriately.

APPENDIX E: AI + IMMIGRATION⁴

Since 2014, Immigration, Refugees, and Citizenship Canada (IRCC) has been in the process of developing a "predictive analytics" system to automate activities currently conducted by immigration officials and to support the evaluation of immigrant and visitor applications. The system, as reported, will or can be used "to identify the merits of an immigration application, spot potential red flags for fraud and weigh all these factors to recommend whether an applicant should be accepted or refused." Public statements from the federal government indicate that the proposed development and adoption of this technology emerged in response to an immigration system encumbered by backlogs and delays. An IRCC analyst confirmed in June 2018 that it is already using some form of automated system to "triage" certain applications into two streams, with "simple" cases being processed and "complex" cases being flagged for review by human counterparts.

While the status of implementing automated decision support systems is not completely clear, it is apparent that at least some decisions are influenced or made by automated review. Concerns have been raised regarding the impact of automated decisions on a variety of human rights: if automated decisions are based on biased data or past decisions, the right to non-discrimination may be violated without human oversight. This is particularly dangerous in the context of immigration, which frequently includes vulnerable populations, limited oversight compared to domestic law, and extremely high impact on wellbeing.

Proponents of automated decisions note that an immigration backlog remains, in part due to a recent surge in asylum seekers, and that long wait times are detrimental to all immigrants. Additionally, while acknowledging that biased decisions can result from automated decisions, they point to the fact that biased decisions can also be made by humans. They also point to examples of good algorithm design that reduced bias in outcomes compared to a human counterfactual.

⁴ Adapted from "Bots at the Gate: A Human Rights Analysis of Automated Decision-Making in Canada's Immigration and Refugee System" © 2018 International Human Rights Program (Faculty of Law, University of Toronto) and the Citizen Lab (Munk School of Global Affairs and Public Policy, University of Toronto). Licensed under the Creative Commons BY-SA 4.0 (Attribution-ShareAlike Licence)

APPENDIX F: IDEAL

Ideal is a talent intelligence application for high-volume recruitment processes that sources, screens, and analyzes candidates in real-time. Using artificial intelligence, Ideal aims to help companies improve the quality of hire, reduce attrition rates, and eliminate recruitment bias.

Ideal's AI software centralizes candidate information gathered from resumes, automated chatbot conversations, and online assessments to help identify the best candidates. Ideal moves beyond the keyword search methods used by other automated hiring systems, and automatically scans, filters, and grades each candidate's resume (as either an A, B, C or D candidate) in real-time.

Ideal assesses candidates based on a combination of disparate data sources such as company performance metrics and past recruitment decisions (e.g. interview invitations, dismissals, employee retention). This enables the system to identify patterns and continuously improve its ability to shortlist strong candidates. Additionally, Ideal optimizes the available talent pool by surveying existing internal and external applicant databases for top candidates. The system also updates past candidate profiles with the latest publicly available information. This removes the need for inconsistent manual screening and allows employers to identify and contact the best candidates in days instead of weeks.

Once candidates have been identified, Ideal's chatbot enables companies to engage with candidates 24/7 and eliminate the need for initial screening calls by asking custom questions like "When are you available to start work?" and "Are you currently enrolled in school or an education program?" Companies are also able to save time by only granting interviews to strong and relevant candidates.

Ideal's data collection and analysis is flexible, and can be programmed to disregard demographic data during collection and analysis in order to prioritize compliance with Canada's employment equity programs. Ideal is currently used by companies such as Indigo and Hot Topic.

Case Study.

Canvas #1: 2018

Step 1 How are different groups experiencing both positive and negative effects?		
? Stakeholders	+ Positive	- Negative
Step 2 What are the potential impacts of this technology?		Step 3 What existing policies and programs are affected?

Case Study.

Canvas #3: Responding Today



<p>Step 1 What are the most important positive and negative implications of this technology today?</p> <div style="margin-top: 20px;"> <p style="text-align: center;">+</p> <p style="text-align: center;">-</p> </div>	<p>Step 2 What are the most important and negative implications of the scenario in 2028?</p> <div style="margin-top: 20px;"> <p style="text-align: center;">+</p> <p style="text-align: center;">-</p> </div>	<p>Step 3 What stakeholder groups experience the strongest gains and the most significant losses.</p> <div style="margin-top: 20px;"> <p style="text-align: center;">+</p> <p style="text-align: center;">-</p> </div>
<p>Step 4 What policy responses today could help to improve outcomes today and prepare for potential futures?</p> <p>Option</p> <p>Benefits</p> <p>Concerns</p>		
<p>Step 5 What are your final policy recommendations?</p> <div style="display: flex; justify-content: space-around; margin-top: 20px;"> <div style="text-align: center;">1</div> <div style="text-align: center;">2</div> <div style="text-align: center;">3</div> </div>		

AI Futures Policy Lab - Vancouver

1. Describe the case study/context - 3 sentences:

2. What are the main opportunities and challenges?

3. What are your top 3 policy recommendations to address these opportunities and/or challenges?