

CIFAR

**CIFAR'S AI & SOCIETY PROGRAM PRESENTS:
A REPORT ON THE
AI FUTURES + HEALTH
POLICY LAB**

HELD IN PARTNERSHIP WITH:
CIFAR's AI & Society Program & BC Ministry of Health

WITH EXPERTISE PROVIDED BY:
Brookfield Institute for Innovation + Entrepreneurship

ABOUT CIFAR

CIFAR is a Canadian-based global charitable organization that convenes extraordinary minds to address the most important questions facing science and humanity.

By supporting long-term interdisciplinary collaboration, CIFAR provides researchers with an unparalleled environment of trust, transparency and knowledge sharing. Our time-tested model inspires new directions of inquiry, accelerates discovery and yields breakthroughs across borders and academic disciplines. Through knowledge mobilization, we are catalysts for change in industry, government and society. CIFAR's community of fellows includes 20 Nobel laureates and more than 400 researchers from 22 countries. In 2017, the Government of Canada appointed CIFAR to develop and lead the Pan-Canadian Artificial Intelligence Strategy, the world's first national AI strategy.

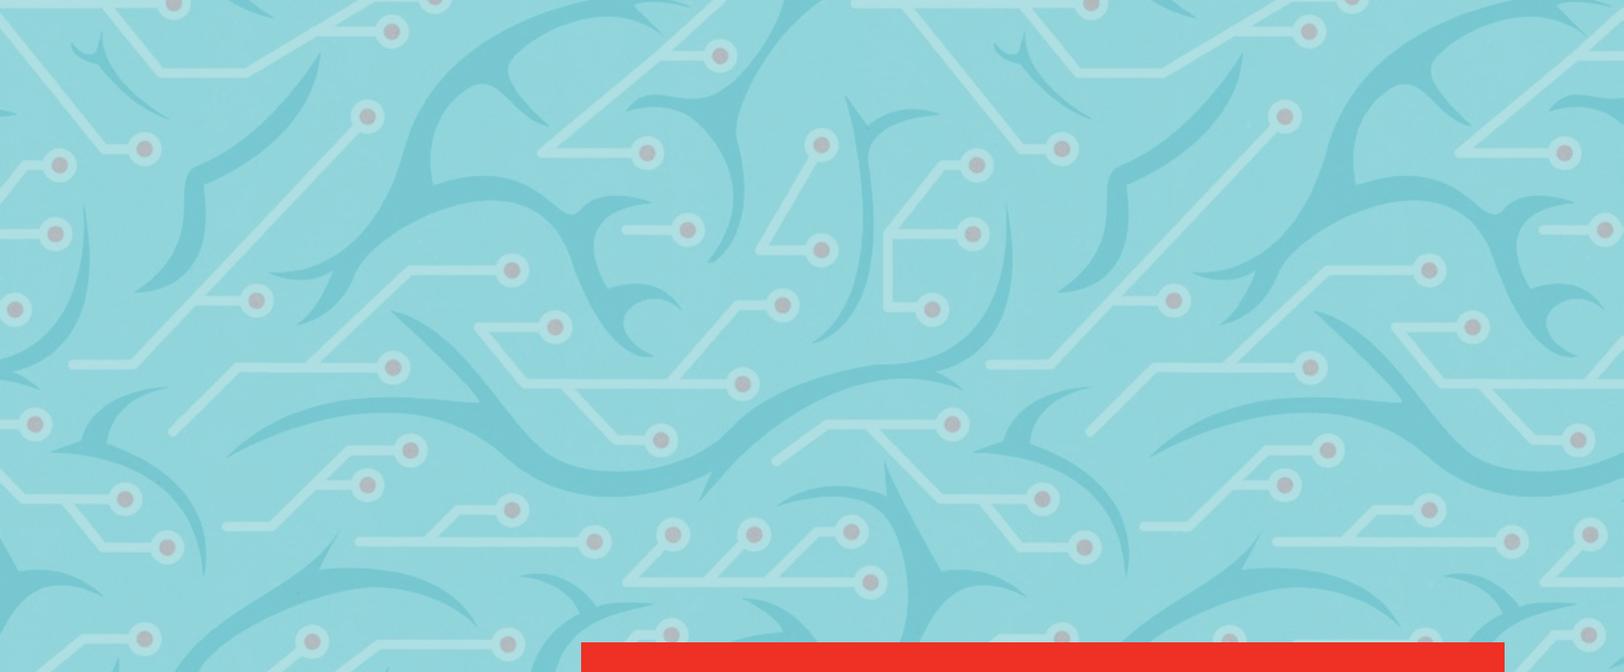
CIFAR is supported by the governments of Canada, British Columbia, Alberta and Quebec, Canadian and international partners, as well as individuals, foundations and corporations.

Copyright © 2020 CIFAR

All rights reserved. No part of this material may be reproduced, distributed, or transmitted in any form or by any means, including screen capturing, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews.

TABLE OF CONTENTS

CIFAR's AI & Society Program	4
Overview	5
Summary of Key Recommendations	5
Case Studies	6
AI Futures Policy Lab Activities	7
Case Studies and Policy Recommendations	9
Key Insights	12
Appendices	14



CIFAR'S AI & SOCIETY PROGRAM

The AI & Society program at CIFAR brings together experts from across sectors, including academia, industry, law, ethics, healthcare, and government, to engage in in-depth discussions on some of the most important and timely challenges that AI is posing to society. These dialogues deliver new ways of thinking about issues and drive positive changes in the development and deployment of successful, beneficial and socially responsible AI technologies.

The AI & Society program is the fourth pillar of the CIFAR Pan-Canadian AI Strategy, a \$125-million investment by the Government of Canada in Canada's leadership in machine learning and training. CIFAR receives additional support for the strategy from Facebook and the RBC Foundation.

OVERVIEW

In 2018, CIFAR and the Brookfield Institute for Innovation + Entrepreneurship (BII+E) launched a series of five workshops to engage policy innovators in conversations about the public policy implications of artificial intelligence (AI). The AI Futures Policy Lab series brought together over 125 policymakers from across Canada to learn about existing and potential AI-enabled capabilities and applications, to explore the policy implications of AI, and to develop policy responses. In each of the workshops, participants developed a variety of policy recommendations to respond to a specific case study.

On February 26, 2020, CIFAR, in partnership with the British Columbia (B.C.) Ministry of Health, held its first AI Futures Policy Lab with a focus on its application to the health sector. The lab, developed with expertise from the BII+E, brought together 21 policymakers, academics, health practitioners, and patient advocates with the aim to:

- Develop a clearer understanding of current AI capabilities;
- Raise awareness of how AI is deployed in real-world health applications
- Discuss the benefits of AI, as well as the social, ethical, economic, and political challenges associated with AI-driven health applications; and
- Identify immediate next steps policymakers need to take in order to support effective deployment and adoption of AI in the health sector in the next five years.

To achieve these aims, the lab featured presentations to provide participants with background information to guide their discussions, engaged participants in hands-on activities that incorporated real-life case studies and canvases, and group feedback sessions. Through these activities, participants gained a deeper understanding of the current AI-enabled capabilities and applications needed to prompt critical thinking surrounding the use of AI in healthcare, facilitate the exploration of potential future scenarios, and develop policy recommendations to inform government responses to the deployment of AI in healthcare.

Participants were divided into groups of four and five, and presented with a case study of an existing AI application in the health care sector. These case studies were curated to reflect the wide-range of AI's applicability within healthcare, and the structural dynamics surrounding its use. Facilitators guided each respective group through their case study and prompted questions to the group to explore the application's impact on stakeholders, current policy, as well as social, economic, and value systems. During the final session, participants were reshuffled so that a member of each case study group could present their recommendations to a new group and receive feedback. The agenda developed for the lab is provided in [Appendix A](#).

SUMMARY OF KEY RECOMMENDATIONS

The group's recommendations shared several overlapping themes:

1. Modernization of data regulation and privacy legislation: Current data regulation and privacy legislation should be revised to address related concerns.
2. Assessment, monitoring and maintenance of AI applications. The ability to evaluate applications, the training data used to develop them, and to establish benchmarks.
3. Importance of running pilot projects and implementing AI-enabled applications in stages. Pilot projects represent a safe and efficient measure for testing the efficacy of new technologies and assessing their impact.

CASE STUDIES

The AI Futures + Health Policy Lab focused on five case studies, highlighting the application of AI on selected areas or elements of the healthcare system. They included public health, public-private partnerships, mental health, health care administration, and radiology treatment. Participants were organized into groups of approximately five people and assigned a relevant case study based on their sector and area of expertise.

PUBLIC HEALTH: FINDER

Existing approaches to identifying suspects of foodborne illness are being augmented with real-time signals at scale, enabling faster detection and more efficient use of public health services, such as health department inspections. FINDER ([Appendix B](#)) is a machine-learning epidemiology system that uses aggregate, anonymized Google search and location information to detect unsafe restaurants, and uses AI to support public health inspections.

PUBLIC PRIVATE PARTNERSHIP: GOOGLE STREAMS

A growing number of public healthcare organizations have entered into partnerships with private companies for the purpose of assisting with the development of tools that could advance patient care and improve outcomes. This case study profiled the partnership between the Royal Free NHS Foundation Trust and Google DeepMind as well as the resulting technology Google Streams ([Appendix C](#)), a mobile medical assistant for clinicians to detect and flag acute kidney injury (AKI) in patients. Using data analytics on a range of patient data, such as blood test results, Google Streams determines whether a patient is at risk of developing AKI and sends an instant alert to clinicians who are able to provide recommendations for preventative action.

MENTAL HEALTH: TREE HOLE

Early detection is one of the key methods to reducing the number of deaths by suicide, and intelligent technology can now help identify individuals at risk. Participants in this group spent their time discussing Tree Hole ([Appendix D](#)), an AI program that applies semantic analysis to monitor and analyze social media posts on the Chinese social media platform Weibo. It identifies individuals at risk of suicide early on, and provides them with access to psychological help.

HEALTH CARE ADMINISTRATION: LKS-CHART

The Li Ka Shing Centre for Healthcare Analytics Research & Training (LKS-CHART) at St. Michael's Hospital in Toronto, Ont., ([Appendix E](#)) applied advanced data analytics to determine the optimal number of Nursing Resource Team (NRT) nurses to hire for the upcoming year, based on historical nurse absence data. These predictions can be used to determine the optimal number of NRT nurses needed, which in turn will minimize overtime and agency expenses.

RADIOLOGY TREATMENT: CORAL REVIEW

Coral Review ([Appendix F](#)), a software solution developed at the University Health Network in Toronto, Ont., is a peer learning tool where radiologists can anonymously review medical imaging diagnoses, improving overall diagnostic accuracy. Coral Review scans through thousands of existing medical images for ones similar to a particular patient and performs diagnosis recommendations to the attending physician. Coral Review users are directed to evaluate the overall trend of image quality for a case, and are prompted to identify whether specific image quality criteria is met throughout the review. The whole process takes less than ten minutes for a reviewer depending on the case complexity.

AI FUTURES POLICY LAB ACTIVITIES

1. SPECTROGRAM

The lab began with an ice-breaker activity known as a spectrogram, in which all participants organize themselves along a continuum in an open space using one wall in the room to represent two views. One side is dedicated to those who “strongly agree” and the opposite wall represents participants that “strongly disagree.” The exercise allows people to have a visual representation of opinions within the group and provides the opportunity for participants to get to know a little about one another. These statements were deliberately designed to be slightly ambiguous, enabling participants to interpret the statements on their own terms.

The prompts for this activity included:

- Karaoke is always fun
- I have a good understanding of AI / I am knowledgeable about AI
- I am optimistic about the potential for AI in health care

2. HEALTH SECTOR INFORMATION, ANALYSIS + REPORTING

Economist Patrick Day, Health Sector Information Analysis and Reporting (HSIAR), provided an overview of the B.C. Ministry of HSIAR division. Day highlighted the potential for AI to create administrative efficiencies, leverage analytical resources, and improve classification and predictive models in HSIAR. Day shared insight into current projects that are being developed within HSIAR. This included:

- Improving audit recoveries from health care billing errors and fraud, as well as the selection of physicians, pharmacies, and patients to audit
- Classifying biological diseases
- Predicting the risk of readmission
- Forecasting health human resources

3. AI 101

Mark Schmidt, Canada CIFAR AI Chair and an associate professor in the Laboratory for Computational Intelligence at the University of British Columbia, provided participants with an overview of key concepts and techniques related to the emergence and current capabilities of AI, with a primary focus on machine learning (ML) using supervised learning techniques. Schmidt provided a number of examples of current ML applications that are widely used today to enhance participants’ understanding. Schmidt concluded his presentation by highlighting areas of concern. This included:

- Autonomous weapons, such as drones equipped with guns and facial recognition technology
- Fake content generation, including the use of machine learning to produce fake audio or video
- Attacks on machine learning systems, including exploiting their vulnerabilities in ways that alter their performance

4. SOCIAL AND ETHICAL CONSIDERATIONS OF AI

Sarah Villeneuve, Policy Analyst at BII+E, provided an overview of the key social and ethical considerations of AI, such as privacy, bias, explainability, and accountability. Her presentation included a number of examples to illustrate why these considerations are important at every step of the design and implementation phases. This included, for example, bias in facial recognition systems, safety flaws in autonomous cars, and the explainability of intelligent diagnostic systems. Villeneuve also provided participants with questions they could use to think more critically about their case studies. These included:

- Where has the data come from? Who collected it, and why? Which groups are not represented in the data?
- What are sufficient ways to measure success in developing safe systems?
- Are these processes in place for end users to seek recourse for errors made, or harms caused, by the system?

5. ROUNDTABLE: ANALYZING CURRENT AI APPLICATIONS IN HEALTH

Within their pre-assigned groups, participants were presented with an example of a current AI application (FINDER, Google Stream, Tree Hole, LKS CHART, and Coral Review). Each participant within the group was given time to read the case study and discuss any preliminary questions. 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 the case study (and whether this is a positive or negative impact); the potential impact 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.

6. EXPLORING THE FUTURE OF AI IN HEALTHCARE

Shawn Gervais, VP, Strategic Foresight, Digital Technology Supercluster, provided insight into how policymakers, healthcare practitioners, and organizations could use foresight to think more critically about current technological developments and sectoral trends. He framed the use of foresight as “preparation”, rather than “prediction,” to ensure participants approached foresight exercises as a way to understand possibilities and different perspectives that could inform a strategy. His presentation included a number of existing AI-driven services and applications that are applied in the realm of healthcare. Shawn also identified some signals of potential future developments in patient experiences and access to healthcare, including:

- The emergence of the citizen physician
- Growth of open source medicine

7. ROUNDTABLE ACTIVITY: EXPLORING AI IN 2025

Facilitators then led their groups into an open discussion about how their AI application may look in 2025. Participants were free to imagine future scenarios and how each application may develop and impact individuals, communities, and policies, as well as social, cultural, political, and economic processes within the next five years.

8. ROUNDTABLE ACTIVITY: 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 canvas and group discussion, and to think about the short- and long-term opportunities and challenges related to the AI application in their case study.

Participants then chose the main opportunity and/or challenge they felt was most appropriate to address or the most interesting, and collaboratively developed a set of policy responses that could assist in achieving that goal. Each participant was provided with a one-page template ([Appendix I](#)) that allowed them to write a short description of the case study they examined, the related opportunities and challenges, the primary goal their group had identified, and the resulting policy recommendations developed.

9. ROUNDTABLE ACTIVITY: RED TEAMING + SHARING BACK

One participant from each group was selected as a representative and placed within a new group with the purpose of giving a short overview of the case study they had explored, the associated opportunities and challenges, primary goal, and related policy recommendations that were developed based on the one-page template they completed in the previous activity. Once they had shared this information, the other participants were required to critique the policy recommendations with the intention of strengthening them.

CASE STUDIES AND POLICY RECOMMENDATIONS¹

PUBLIC HEALTH: FINDER

Throughout the day, participants in this group identified a number of opportunities and challenges related to the use of FINDER. On one hand, they believed FINDER provided the ability to identify unsafe restaurants more efficiently, and optimize food-related inspection systems. They also saw FINDER as a potential incentive for restaurant owners to better manage their reputation and increase their accountability of business and staff practices. However, there were a number of concerns related to the application, particularly associated with misuse, surveillance, and economic consequences. Participants recognized the potential economic implications this application could have if restaurant owners and others were aware of its use, and maliciously triggered the application's response. There was discussion around the potential to disrupt the current random sampling health inspection model and how this might bias the food inspection system. Participants were also concerned about the level of public surveillance enabled by the app, and what other forms of surveillance it could encourage. Participants questioned whether there was another way to identify suspect restaurants that didn't involve the use of personal data (such as an individual's location and search queries). The use of search query data also caused participants to question whether the data was not inclusive of non-english speakers.

This group identified the main goal of the app was to eliminate food-related illness and avoid unintended consequences, such as false positives which could negatively affect businesses. This involved ensuring the accuracy and efficiency of the method used to pin-point suspect restaurants.

Group recommendations:

- To develop an algorithmic impact assessment for health applications, including an assessment of the training data set for greater transparency;
- Conduct annual monitoring and reassessment after the application has been implemented; and,
- To take a staged approach to implementation, starting at the neighborhood level, before implementing the application at the municipal, provincial or national level.

During the red-teaming activity, feedback from other members of the lab encouraged this group to consider the social media implications and possibilities of sabotage or push from competitors. For example, they explored the question of whether individuals generate false positives for businesses they were in competition with or disliked by making searches that would deliberately link to their location history. The red-teaming feedback encouraged the group to consider whether the benefits of this technology outweighed the risks. The group questioned whether this system was valuable for implementation as a public health resource. Feedback also encouraged participants to consider the importance of public consultation before implementing this kind of tool.

PUBLIC PRIVATE PARTNERSHIP: GOOGLE STREAMS

This case study prompted participants to discuss the tension between the beneficial impacts of new technologies and resistance to change within the healthcare system. The benefits identified included the use of AI to assist in health screening and improve patient outcomes. They see the potential for this application to reduce the burden on physicians by enabling patients and physicians to cooperatively diagnose AKI in real-time. The use of this application is also expected to decrease the cost of detecting AKI, per patient. Participants were optimistic that this kind of application could be scaled to help identify other diseases.

However, the primary challenge this group identified was the healthcare systems resistance to change. This could reduce the level of adoption and integration of beneficial technology into the healthcare system. The extent to which this application can also pose challenges to physicians, by altering or augmenting their roles and responsibilities was also raised as a potential issue. Another barrier participants recognized was establishing a benchmark for the tolerance for error this technology could have within the healthcare context. This is strongly related to the levels of citizen trust. Data privacy and security was a significant concern within this group, due to the nature of public private partnerships. They also discussed the need for the long-term measures of equity, to demonstrate how the benefits of a particular technology and partnership are spread across the sector and target populations.

¹ Disclaimer: The following policy recommendations were developed by participants through an exercise designed to have participants explore existing policy levers in relation to specific case studies. These do not represent the views of CIFAR, BII+E, or the BC Ministry of Health.

This group identified a dual goal of using AI to increase life expectancy while ensuring the technology used is stable, robust, trustworthy, secure, affordable and equitable. This includes building public private partnerships that prioritize data security, privacy, public consent, and follow established principles and policy.

Group recommendations

- Encourage cooperative development, adoption, and integration of AI technologies through legislation, and thus promoting greater collaboration between siloed public bodies. This can be achieved by developing provincial technology strategies (like [Canada's Digital Charter](#)) that establish procurement principles and policies. These strategies should prioritize data security, privacy, and public consent.
- Incorporate perspectives and considerations from Indigenous communities for more inclusive data governance. This could include the [Declaration on the Rights of Indigenous Peoples Act](#) (DRIPA) and [Ownership Control Access and Possession](#) (OCAP).
- Modernize privacy legislation, potentially informed by [The Five Safes](#) framework.
- Conduct maintenance and quality assurance of machine learning algorithms.
- Develop policy that enables and supports recourse when an algorithm results in a mis-diagnosis and/or patient harm.

Feedback that was provided during the red-team discussion encouraged the group to emphasize the need for consent. They also highlighted that having clear data governance is critical for public-private partnerships, particularly as a mechanism for control and for increasing public trust. Questions were raised around how to involve the public in the development and governance of these partnerships, and prompting consideration of the need for public education and awareness raising.

MENTAL HEALTH: TREE HOLE

In their discussions about Tree Hole, participants highlighted a number of opportunities this application has for individuals who are struggling with their mental health, as well as their communities and mental health care practitioners. Participants saw the potential for Tree Hole to enhance the quality of life for individuals at risk of suicide by preventing suicides through more efficient detection and increased monitoring. They were also optimistic about the improvements to the health delivery. Importantly, they pointed to the ability for this application to decrease inequalities in access to mental health and health care services while lowering costs of service provision, improving access for communities more broadly. However, participants

also identified a number of areas of concern. They were particularly concerned with the impact this type of application could have on individual autonomy, particularly on people who are not seeking out mental health services. There was a risk that this app could increase stigma surrounding mental health, especially if people were using social media as a way to talk about their struggles due to existing stigma against mental health in their community. Participants in this group were also concerned about individual privacy, since many people may not know they are being “watched” and are being flagged by this application. They also raised questions regarding how this application can adapt to evolving language and cultural references or contexts.

This group's main recommendation was to develop and launch a pilot project to test this application in B.C. They saw this as a way to further advance the application to increase the benefits for identifying at-risk individuals, while also detecting areas where the app could be improved in order to mitigate privacy and security challenges.

Group recommendations

- Identify and convene relevant stakeholders to ensure the pilot project is supported by expert insight and the lived experiences of patients and health practitioners.
- Provide funding for pilot projects and initiatives that seek to test existing AI-driven health care applications that have been proven to provide benefit, in order to ensure they work in different contexts.
- Further develop existing policy and legislation surrounding data regulation to define the limits of data that will be collected through this application.
- Develop infrastructure for implementation and knowledge sharing for wider use.

The red-teaming activity raised a number of questions and challenges associated with the use and effectiveness of this application. One point raised was that Tree Hole does not address the root problem of suicidality. For this reason, they believed this application would provide the most value if it were implemented without the consent and knowledge of those who were being monitored, however participants saw this as unethical. Additionally, participants felt as though individuals could discover how to game the system and jump the queue by publishing posts that would categorize them as high-risk. Instead of being applied to social media, the red-team participants wondered whether there was an opportunity to implement this within existing care systems. There was also unclarity surrounding the cost of implementing and maintaining the application.

HEALTH CARE ADMINISTRATION: LKS-CHART

Throughout their exploration of LKS-CHART, participants thought about how this kind of application could be applied within the health care system or in other public services. They believed this tool provided greater system efficiency by introducing more consistent scheduling for nurses and balancing overtime costs against patient demand. Throughout these discussions, a number of concerns were also identified, particularly surrounding the challenges related to adoption, such as system resistance or lack of buy-in. The group raised potential issues surrounding the quality of data used to inform this system, such as the inventory of existing data or whether relevant data even exists. Participants also expected implementation challenges due to union concerns and collective agreements of nurses and other health care workers. They were concerned that a focus on cost optimization could deprioritize or risk patient health care. They also cited the app's lack of monitoring and evaluation mechanisms for assessing healthcare outcomes as a challenge, as well as matching the appropriate skills and expertise with patient demand.

The group identified the primary goal of the application was to apply the scheduling tool to healthcare operations more broadly.

Group recommendations

- Create government-led incentives for adoption and buy-in to tackle privacy concerns.
- Adopt incrementally to avoid push back, and gradually scale up across the province.
- Develop and implement a secure data storage infrastructure.
- Engage stakeholders early (to understand preferences and early buy-in).
- Ensure apps are developed based on needs of practitioners.
- Identify and manage HR implications associated with data privacy and security.
- Develop a better understanding of the dynamics of patient demand in order to efficiently predict and balance staff required.

During the red-teaming, participants raised several points for consideration for organizations that choose to use LKS-CHART. The first consideration was related to the quality of data, specifically calling for organizations to evaluate whether they already have the right kind of data, and enough of it, to make accurate decisions related to their staff. Additionally, participants suggested that organizations ensure they have the ability to add data into the system on an ongoing basis in order to adapt to changes in circumstances, such as a

disease outbreak. The second consideration concerned the infrastructure needed to implement this application. Participants were aware that some hospitals may not have the right technical infrastructure in place to adopt this properly. Thirdly, participants raised the need to establish measures to evaluate and monitor adoption to ensure it was delivering the intended impact. This requires establishing a baseline.

RADIOLOGY TREATMENT: CORAL REVIEW

Within this group, participants discussed the potential for Coral Review to provide better patient outcomes, faster treatment, increase patient trust, reduce clinical workload, and cost savings. However, participants identified a number of concerns. For example, they were wary that the use of this system could foster favouritism towards AI-driven decisions over physician decisions. They also recognized that the use of this system may result in job loss or require some occupations to retrain or reorient their skillsets. Additionally, they highlighted that this system could create downstream backlogs, and acknowledged there would need to be checks and balances in place to mitigate these potential challenges.

This group decided their primary goal would be to generalize Coral Review's scope across all areas of radiology therapy.

Group recommendations

- Develop a standardized regulatory framework.
- To scale the application across other disease areas.
- To expand the use of the app to other clinical centres.
- Create occupational training for clinicians to effectively use AI-driven health applications.
- Create funding incentives to support the adoption of Coral Review.
- Develop cross sector teams working across the country to tackle the changes.

Through the red teaming activity, feedback from the other participants questioned whether the application's preferred rate is high enough to warrant the use of this system over traditional methods. Before implementing this system more widely, it was suggested to include patient outcomes in the training data for this system.



KEY INSIGHTS

While each group was provided with a different AI application to explore throughout the day, discussions and final recommendations shared several overlapping themes worth highlighting.

1. Modernization of data regulation and privacy legislation

All groups underlined that current data regulation and privacy legislation may need to be revised in order to properly address concerns related to new challenges associated with privacy, surveillance, and data collection, use and sharing practices.

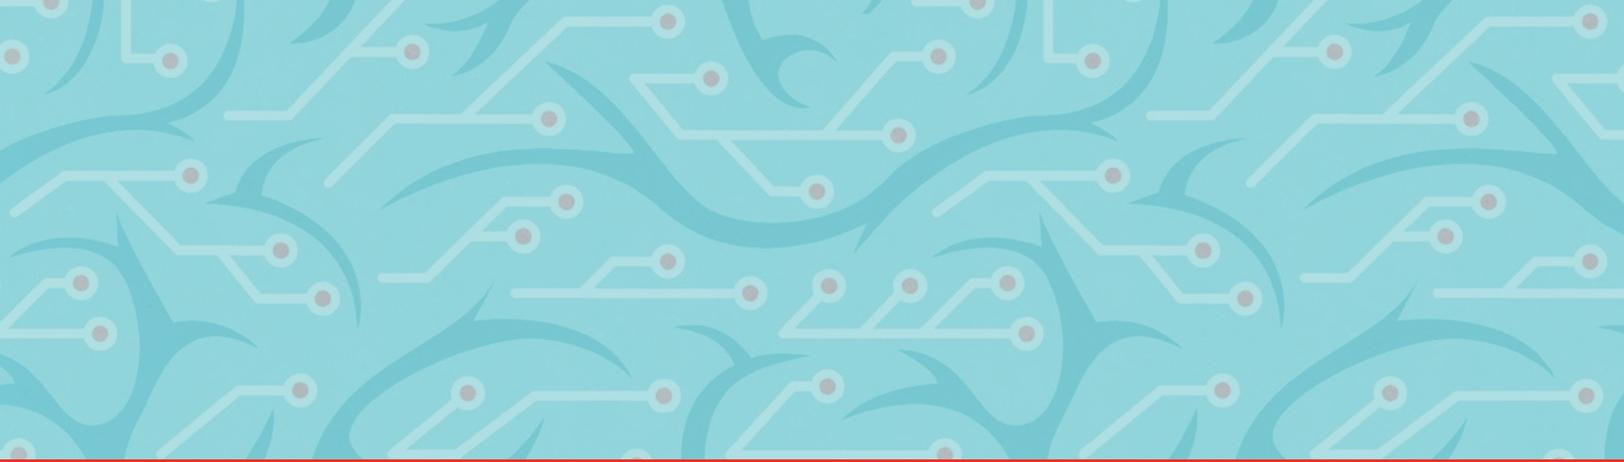
2. Assessment, monitoring, and maintenance of AI applications

Some groups discussed the development and use of Algorithmic Impact Assessments to evaluate each of the applications technical operations, as well as the training data that was used to develop them. Participants expressed the need to establish a benchmark for the tolerance of error to measure any unintended consequences of the applications' use, and ensure there is an established avenue for recourse if harms occur. Several participants emphasized the need to incorporate Indigenous data governance into evaluation and monitoring frameworks. Ideally, an assessment and monitoring framework would enable organizations to measure how the benefits of applications are being distributed across sectors, and how they target populations, in order to address any potential disparities in equity of impact.

3. Importance of running pilot projects and implementing AI-enabled applications in stages

These measures would help organizations test the effectiveness and assess the impacts of new AI-enabled applications when implemented in a new context. Participants stressed that pilot projects and staged implementation processes should be informed by experts and the lived experiences of the target population.

These three themes point to important considerations that are likely needed to be taken when exploring the potential use and implications of an AI-driven application, both in the health care context and beyond.



NEXT STEPS

CIFAR will launch a virtual AI Futures Policy Lab tool kit on cifar.ca in summer 2020. Composed of printable worksheets, a comprehensive facilitation guide, and complementary informational videos showcasing leading Canadian AI experts, this online resource will enable groups and organizations to run their own AI Futures Policy Lab. For more information, please contact Gaga Boskovic (gaga.boskovic@cifar.ca).

ACKNOWLEDGEMENTS

This CIFAR AI Futures Policy Lab was held in partnership with the British Columbia Ministry of Health with expertise and report writing provided by the Brookfield Institute for Innovation + Entrepreneurship (BII+E), as part of CIFAR's AI & Society program.

Lead author: Sarah Villeneuve, Policy Analyst,
Brookfield Institute for Innovation + Entrepreneurship



Ministry of
Health



APPENDICES

APPENDIX A: AGENDA

AI FUTURES + HEALTH POLICY LAB

At the end of this workshop, participants will have:

- developed a clearer understanding of current capabilities of AI;
- awareness of how AI is being used in the health sector, based on real-world applications;
- discussed benefits as well as social, ethical, economic, and political challenges associated with AI-driven health applications;
- identified immediate next steps policymakers need to take in order to support effective deployment and adoption of AI in the health sector in the next five years.

Background reading materials (optional):

- Brookfield Institute: [Intro to AI for Policymakers](#)
- Alliance for AI in Healthcare: [Artificial Intelligence in Healthcare: A Technical Introduction](#)
- CIFAR: [Rebooting Regulation: Exploring the Future of AI Policy in Canada](#)

TUESDAY FEBRUARY 25, 2020

6:00 p.m.	Pre workshop Dinner Opening Remarks, Rebecca Finlay, CIFAR Overview of provincial AI for Health landscape, Heather Davidson, Ministry of Health Overview of Pan-Canadian AI Strategy, AI for Health Task Force, Elissa Strome, CIFAR
-----------	--

WEDNESDAY FEBRUARY 26, 2020

8:30 – 9:00 a.m.	Light Breakfast & Networking
9:00 – 9:15 a.m.	Opening Remarks: Amy Cook, CIFAR; Martin Wright, HSIAR
9:15 – 9:30 a.m.	Warm-Up: Spectrogram
9:30 – 9:50 a.m.	Health Sector Information, Analysis + Reporting: Patrick Day, HSIAR
9:50 – 10:35 a.m.	AI 101: Mark Schmidt, Canada CIFAR AI Chair in B.C.
10:35 – 10:55 a.m.	Break
10:55 – 11:20 a.m.	Social and Ethical Considerations of AI: Sarah Villeneuve, BII+E
11:20 a.m. – 12:20 p.m.	Roundtable Activity: Analyzing Current AI Applications
12:20 – 1:20 p.m.	Lunch
1:20 – 1:40 p.m.	Exploring the Future of AI: Shawn Gervais, VP, Strategic Foresight, Digital Supercluster
1:40 – 2:10 p.m.	Roundtable Activity: Exploring AI in 2025
2:10 – 3:10 p.m.	Table Activity: Taking Action Today
3:10 – 3:25 p.m.	Break
3:25 – 4:20 p.m.	Roundtable Activity: Red Teaming
4:20 – 4:30 p.m.	Closing Remarks and Next Steps: Amy Cook, CIFAR

APPENDIX B: MACHINE-LEARNING EPIDEMIOLOGY USING FINDER

AI Application: Machine Learning

Health Focus Area: Public Health

The Foodborne Illness Detector in Real time, also known as FINDER, is a machine-learning epidemiology system able to detect unsafe restaurants. This model was developed by a team of Google researchers and public health experts.² FINDER uses aggregate, anonymized Google search and location information, collected from users who have enabled location services on their smartphone device for the use of Google maps, or other location-based apps. The model first works to identify users who have searched terms such as “vomiting” or other symptoms that are indicative of food poisoning. It then investigates their location history to determine which restaurants the user had been to recently. For each applicable restaurant, the system then calculates the proportion of users who visited it and later searched for symptoms related to foodborne illnesses.² FINDER then flags suspect restaurants for health department officials to inspect.³

When deployed in Las Vegas and Chicago, FINDER was 3.1 times as likely to identify “unsafe restaurants compared to current health inspection methods.”³ Additionally, researchers found that 52 per cent of restaurants flagged by FINDER were deemed unsafe upon inspection.³ In fact, FINDER is more likely to identify problematic restaurants than routine inspections. Comparatively, just 25 per cent of restaurants that undergo routine inspections are deemed unsafe.³ FINDER also seems to be a stronger indicator of foodborne illness than customer complaints, which were found to be accurate 39 per cent of the time. This is primarily because people often assume they received food poisoning from the last place they ate, but according to FINDER results, this is only accurate in 61 per cent of the cases identified.³ Overall, FINDER provides a complementary model to existing approaches for identifying suspects of foodborne illness by providing new, real-time signals at scale.

APPENDIX C: GOOGLE STREAMS

AI Application: Machine Learning

Health Focus Area: Public-Private Partnership

In partnership with Royal Free NHS Foundation Trust, Google DeepMind uses Google Streams, a mobile medical assistant for clinicians, to help detect and flag acute kidney injury (AKI) in patients, a condition where a patient’s kidney suddenly stops working properly. Using data analytics on a range of patient data, such as blood test results, Google Streams determines whether a patient is at risk of developing AKI and sends an instant alert to clinicians who are able to provide recommendations for preventative action. The data is stored on a mobile app, which conducts the analyses on behalf of the clinicians, reducing the administrative burden on the staff.

The Royal Free London NHS Foundation Trust controls all of the data that Google Health accesses. Google Health acts only as a data processor, accessing personal data only for the purpose of providing Streams. Google Health is not permitted to use personal data for any other purpose. Streams is fed with information from the other information technology systems used by the hospitals. The personal data in Streams is subject to a ten-year retention period.

[Royal Free NHS Foundation Trust](#)

[Business Insider](#)

[Google DeepMind](#)

² Sadilek, Adam, Stephanie Caty, Lauren DiPrete, Raed Mansour, Tom Schenk, Mark Bergtholdt, Ashish Jha, Prem Ramaswami, and Evgeniy Gabrilovich. ‘Machine-Learned Epidemiology: Real-Time Detection of Foodborne Illness at Scale’. *Npj Digital Medicine* 1, no. 1 (6 November 2018): 1-7. <https://doi.org/10.1038/s41746-018-0045-1>.

³ Hu, Jane C. ‘Your Google Data Could Pinpoint Where You Got Food Poisoning’. *Quartz*. Accessed 1 April 2020. <https://qz.com/1452799/your-google-data-could-pinpoint-where-you-got-food-poisoning/>.

APPENDIX D: TREE HOLE

AI Application: Natural Language Processing Health Focus Area: Mental Health

According to the World Health Organization, approximately 800,000 people die by suicide every year, and it is the second leading cause of death among 15–29 year olds globally.⁴ However, suicides are preventable — a number of measures can be taken at individual, sub-population, and population levels to prevent suicide and suicide attempts. Early detection is one of the key methods to reducing the number of deaths by suicide.

Tree Hole, launched in 2018, is an AI program that applies semantic analysis programming to monitor and analyze social media posts on the Chinese platform Weibo. It aims to identify individuals at risk of suicide early on, and to provide them with access to psychological help.⁵ This largely benefits individuals who do not have proper access to professional mental health services, such as those who live in small or rural areas. In these cases, Tree Hole can play an important role in early detection and help to connect those in need with resources.

The program automatically scans Weibo's platform every four hours, using a knowledge graph of suicide notions and concepts to identify posts that contain words and phrases that may be indicative of suicidal tendencies.⁶ When concerning posts are found, Tree Hole's system notifies a network of 600 people, made up of psychologists, consultants, and volunteers.⁷ These psychologists, consultants, and volunteers reach out to individuals who post about their mental health online.

Tree Hole automatically classifies social media posts in ten levels, depending on the perceived severity.⁸ Posts that are categorized at the highest level, ten, contain details such as suicide method, date, and time. These are deemed to require urgent action by the Tree Hole team.⁹ Posts that are categorized at level nine indicate a strong belief that a suicide attempt may be made soon.¹⁰ In these cases, an attempt is made by Tree Hole's network to make contact with the individual who posted. If Tree Hole can not make contact, they alert the authorities.¹¹ Posts categorized in below level six and are not deemed to require intervention, as they only include negative words. Since its launch, Tree Hole has reported to have prevented more than 1,000 suicides in China.¹² The application is claimed to be 82 per cent accurate in identifying and classifying suicidal words and phrases.¹³

Applying AI to social media in order to detect and prevent suicide is not new. A number of social media websites have attempted to use similar applications for this purpose. In 2017, Facebook attempted to apply machine learning to identify users who might be at risk of suicide by analyzing words and phrases in posts.¹⁴ However, Facebook was criticized for not obtaining affirmative consent to share information with law enforcement.¹⁵

⁴ 'Suicide'. Accessed 1 April 2020. <https://www.who.int/news-room/fact-sheets/detail/suicide>.

⁵ 'This AI Bot Tracks Suicidal Posts on Social Media to Help Psychologists Save Lives - VICE'. Accessed 1 April 2020. https://www.vice.com/en_in/article/d3ab5q/this-ai-bot-tracks-suicidal-posts-on-social-media-to-help-psychologists-save-lives.

⁶ South China Morning Post. 'AI Bot Finds Suicidal Voices on Chinese Social Media, Helping Save Lives', 18 November 2019. <https://www.scmp.com/tech/apps-social/article/3037917/ai-bot-finds-suicidal-messages-chinas-weibo-helping-volunteer>.

⁷ 'This AI Bot Tracks Suicidal Posts on Social Media to Help Psychologists Save Lives - VICE'. Accessed 1 April 2020. https://www.vice.com/en_in/article/d3ab5q/this-ai-bot-tracks-suicidal-posts-on-social-media-to-help-psychologists-save-lives.

⁸ Wang, Yitsing. 'The Chinese Suicides Prevented by AI from Afar'. *BBC News*, 9 November 2019, sec. Technology. <https://www.bbc.com/news/technology-50314819>.

⁹ 'This AI Bot Tracks Suicidal Posts on Social Media to Help Psychologists Save Lives - VICE'. Accessed 1 April 2020. https://www.vice.com/en_in/article/d3ab5q/this-ai-bot-tracks-suicidal-posts-on-social-media-to-help-psychologists-save-lives.

¹⁰ Wang, Yitsing. 'The Chinese Suicides Prevented by AI from Afar'. *BBC News*, 9 November 2019, sec. Technology. <https://www.bbc.com/news/technology-50314819>.

¹¹ Danni, Fu. 'Keeping an Ear to Weibo's Suicidal Whispers'. *Sixth Tone*, 6 June 2019. <https://www.sixthtone.com/news/1004104/keeping-an-ear-to-weibos-suicidal-whispers>.

¹² South China Morning Post. 'AI Bot Finds Suicidal Voices on Chinese Social Media, Helping Save Lives', 18 November 2019. <https://www.scmp.com/tech/apps-social/article/3037917/ai-bot-finds-suicidal-messages-chinas-weibo-helping-volunteer>.

¹³ South China Morning Post. 'AI Bot Finds Suicidal Voices on Chinese Social Media, Helping Save Lives', 18 November 2019. <https://www.scmp.com/tech/apps-social/article/3037917/ai-bot-finds-suicidal-messages-chinas-weibo-helping-volunteer>.

¹⁴ About Facebook. 'How Facebook AI Helps Suicide Prevention', 10 September 2018. <https://about.fb.com/news/2018/09/inside-feed-suicide-prevention-and-ai/>.

¹⁵ 'Inside Facebook's Suicide Algorithm - Business Insider'. Accessed 1 April 2020. <https://www.businessinsider.com/facebook-is-using-ai-to-try-to-predict-if-youre-suicidal-2018-12>.

APPENDIX E: LKS-CHART

AI Application: Machine Learning

Health Focus Area: healthcare administration

St. Michael's Hospital recently created a Nursing Resource Team (NRT), a team of staff nurses whose role is to cover vacations, sick leaves and other short-term absences on the hospital's nursing units. It is important to determine the optimal size of the NRT. If there are not enough NRT nurses to cover absences, then the hospital relies on staff overtime and external nursing agencies to fill the gaps, which has significant costs. If the NRT is overstaffed, then the costs of the program become unnecessarily inflated due to supply exceeding demand.

LKS-CHART applied advanced data analytics to determine the optimal number of NRT nurses to hire for the upcoming year. Using historical nurse absence data and IBM's PureData for Analytics system, they built a forecast model to predict the weekly number of absences up to a year in advance. These predictions can then be used to optimally determine the number of NRT nurses that will minimize the total of NRT, overtime, and agency costs.

Source: <https://www.chartdatascience.ca/nrt-forecasting-optimization>

APPENDIX F: CORAL REVIEW

AI Application: Machine Learning

Health Focus Area: Radiation Therapy

The accuracy of a patient's diagnosis is dependent largely on the quality of the image that the technologist produces and the interpretation of the images by the radiologist. If medical images are of poor quality, this will reduce the accuracy of the diagnosis. Moreover, physicians need to be available to provide a second opinion to review diagnoses for accuracy, however, their availability may be limited.

Coral Review, a software solution developed at the University Health Network in Toronto, Ont., is a peer learning tool where radiologists can anonymously peer review medical imaging diagnoses, improving diagnostic accuracy. Coral Review scans through thousands of existing medical images for ones similar to a patient's and recommends a diagnosis to the attending physician. Coral Review users are directed to evaluate the overall trend of image quality for a case, and are prompted to identify whether specific image quality criteria is met throughout the review. The whole process takes less than ten minutes for a reviewer depending on the case complexity.

Approximately 2–5 per cent of Joint Department of Medical Imaging patient cases will be reviewed using this software, and in its first month, more than 1,000 images were reviewed – translating to numerous learning opportunities for image quality improvement. Coral Review software is used by several health centres in Ontario, including eight organizations and 19 sites, and JDMI's vision is to eventually roll the solution out across the province enabling system-wide improvement.

APPENDIX G: CANVAS 1

Case Study: Today  **CIFAR**
for innovation + entrepreneurship

How is this affecting people?

<p>⊕ Stakeholders</p>	<p>⊕ Positive</p>	<p>⊖ Negative</p>
-----------------------	-------------------	-------------------

How is this affecting the world?

<p>Social <small>Behaviour, demographics, health + leisure</small></p>	<p>Technological <small>Research, innovation + invention</small></p>	<p>Environmental <small>Air, water, pollution, cities + towns</small></p>	<p>Economic <small>GDP, employment, incentives, income + distribution</small></p>	<p>Political <small>Institutions, elections, lobbying, influence + power</small></p>	<p>Values <small>Beliefs, ethics, + priorities</small></p>
---	---	--	--	---	---

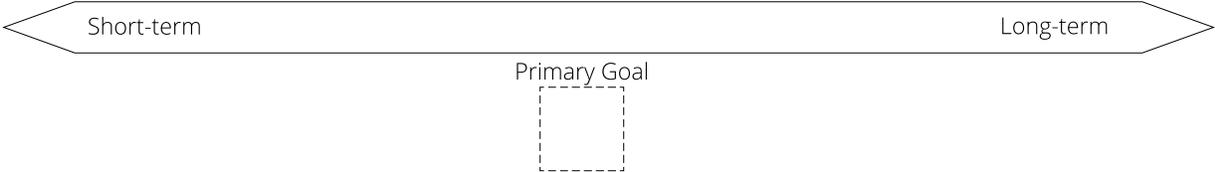
APPENDIX H: CANVAS 2

Case Study:

Taking Action



What are the goals that we are trying to achieve?



What policy responses could help to achieve that goal?

Option

Benefits

Concerns

APPENDIX I: ONE-PAGE TEMPLATE

AI Futures Policy Lab



1. Describe the case study/context in ~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?

CIFAR

CIFAR is a Canadian-based global charitable organization that convenes extraordinary minds to address the most important questions facing science and humanity.

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

cifar.ca/ai

 [@CIFAR_News](https://twitter.com/CIFAR_News)
 Facebook.com/CIFAR
 Linkedin.com/company/CIFAR

For partnership opportunities, contact donations@cifar.ca