AI Futures Policy Lab:
Edmonton

October 2018

brookfield institute
for innovation + entrepreneurship

CIFAR

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).

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+E

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.

/CIFAR

@CIFAR_News

CIFAR

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

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.

/BrookfieldIIE

@BrookfieldIIE

The Brookfield Institute for Innovation + Entrepreneurship

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

TABLE OF CONTENTS

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 to generate greater awareness of the long-term implications of AI and build capacity to develop agile AI policy in Canada. The first lab took place on June 25, 2018 in Toronto, with participation from 18 emerging policy leaders.

On September 20, 2018, CIFAR and BII+E hosted the second AI Futures Policy Lab in Edmonton, Alberta. This event brought together 23 emerging policy leaders with the aim to:

+ Build capacity of future public service leaders to understand the policy implications of AI and respond appropriately;
+ Provide policymakers with a direct line of sight into the AI sector: the myths and hype, the evolving state of technological advances, and potential applications;
+ Contribute 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 AI Futures Policy Lab Pilot in Toronto was integrated into the design of the Edmonton lab's agenda and content. Rather than focusing on just one potential future scenario throughout the day, participants analyzed a pair of case studies associated with a specific policy domain (i.e. housing, justice, education, and health), one current state example and one potential future scenario. At the end of the day, groups produced and presented policy recommendations related to the two case studies they analyzed throughout the afternoon. The agenda developed for the day is provided in appendix A.

CASE STUDY POLICY DOMAINS

Prior to the lab, four sets of case studies were developed. Each set was associated with a specific domain: housing, legal, education, or health. Participants were organized into groups of 4-5 and assigned 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. They were also presented with a scenario that imagined a future in which a smart-home contractor approached the local municipal government with the proposal of building affordable housing in exchange for the collection of resident data (appendix C).

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 D), an artificially intelligent legal research tool that applies natural language processing to increase lawyer's ability to identify relevant information. This group was then presented with a future scenario in which legal decisions for minor infractions were made by an artificially intelligent agent (appendix E).

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 F), an artificially intelligence class assistant that uses machine learning algorithms and advanced facial recognition to analyze the attention of students listening to online lectures. This group then examined a future scenario in which intelligent devices and toys are integrated within classrooms to help monitor children's ability to learn, track their progress, and optimize their experiences (appendix G).

HEALTH

Advancements in AI capabilities hold enormous opportunities for delivering more efficient health care services in areas such as diagnosis, health monitoring, and treatments. However, this also raises challenges related to patient privacy and discrimination. Participants within this group explored InnerEye (appendix H), a research initiative led by Microsoft that applies computer vision and machine learning algorithms to automatically analyze three-dimensional medical CT (computer tomography) and MR (magnetic resonance) images to identify tumours. The group then explored a future scenario in which an individual's health was constantly monitored by a set of ubiquitous devices and applications that tracked variables such as activity, sleep, speech patterns, expressions, movements, and pulse (appendix I) to proactively diagnose conditions and recommend treatment.

POLICY LAB ACTIVITIES

1. DECOLONIZING SCIENCE + TECHNOLOGY

The day kicked off with a keynote speech by Dr. Kim TallBear, Canadian Research Chair in Indigenous Peoples, Technoscience, and Environment and Associate Professor in the Faculty of Native Studies at the University of Alberta. Dr. TallBear provided a historical analysis of the objectification of Indigenous peoples in western science and technology discourse. Widening the scope of contemporary inclusion and diversity rhetoric to include indigenous peoples, Dr. TallBear highlighted the need for space within the AI R&D community for Indigenous peoples to participate as collaborators, scientists, developers, and regulators. This session also included a Q+A for participants.

2. THE AI THING FROM THE FUTURE

Following Dr. TallBear's talk, each table of participants played two rounds of *The 'AI' Thing from the Future*[1] accompanied by a facilitator. The purpose of this activity was to encourage participants to be creative and think beyond the current reality. Each group was then given five cards, each containing a different prompt: *ARC*, to signify what type of future; *terrain*, defining the thematic context or location of the object; *object*, specifying the type of artifact you are focusing on; *mood*, suggesting how you might feel when experiencing this thing; *and AI*, indicating the technological capability or application that needs to be integrated in your future "thing" (appendix J).

Participants were provided with a template (appendix K) to record their ideas. Each participant was asked to individually imagine a future object, or 'thing', utilizing all card prompts. This was followed by an opportunity to share these ideas with the rest of the group.

3. AI 101

Cam Linke, Interim Executive Director of the Alberta Machine Intelligence Institute (AMII), provided an overview of key terms within the AI field, learning techniques, neural networks, recent achievements and current research areas. Participants were given the opportunity to ask clarifying questions. This presentation enabled participants to enter into subsequent discussions with a better awareness of current technical capabilities and applications.

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

4. ANALYZING CURRENT AI APPLICATIONS

Within their pre-assigned groups, participants were presented with an example of a current AI application (Naborly, ROSS Intelligence, Nestor and InnerEye). Each group was given time to read the case study and discuss any preliminary questions they had before moving on to the canvas. Once the group was comfortable in their understanding of the case study, they turned their focus to the first canvas (appendix L). This canvas prompted participants to think about who is impacted and how (positively and/or negatively); what are the potential impacts of this technology at the local, national, and global levels; and what existing policies and programs does this technology impact? Participants were encouraged to actively contribute by writing their ideas on sticky notes and placing them on the canvas, first individually and then as a group.

5. EXAMINING AI IN 2028

This session was designed to enable participants to analyze a possible future scenario that builds upon their current state case study within the same policy domain. Presented with the future scenario, participants were asked to discuss their concerns or excitement regarding the technology. Similar to the first canvas, the canvas for this session (appendix M) prompted participants to imagine which individuals or groups would be impacted by this technology and how (positively and/or negatively), as well as the impacts this technology would have at the local, national, and global levels. The canvas also prompted participants to imagine alternative future scenarios using the current AI application as a starting point.

6. TAKING ACTION TODAY

During the final session of the day, participants were encouraged to review the two previous canvases and reflect upon their discussions (appendix N). This session asked participants to highlight the most important positive and negative socio-political impacts of both the current and future case studies, as well as which individuals or groups experience the most significant impacts (positive and/or negative).

Participants then collaboratively brainstormed and discussed a range of policy options that would enhance the opportunities afforded by the technology they had analyzed throughout the day, while mitigating the risks it could have on society. This exercise was designed to help participants think through which available policy levers could enable them to respond effectively to the issues at hand. Each group then developed two to three policy recommendations to demonstrate their thinking.

Groups were given 5 minutes to outline their case studies, the associated risks and opportunities, and their policy recommendations to the other groups. Policy recommendations are described in the following section.

POLICY RECOMMENDATIONS

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

DOMAIN: HOUSING

Participants discussed a range of possibilities about how AI could drive change in new and existing communities, and even challenge our conception of communities. With a primary focus on the well-being of residents, the group proposed a model of policy development that relies on a feedback loop of values, testing and revision regarding AI policy:

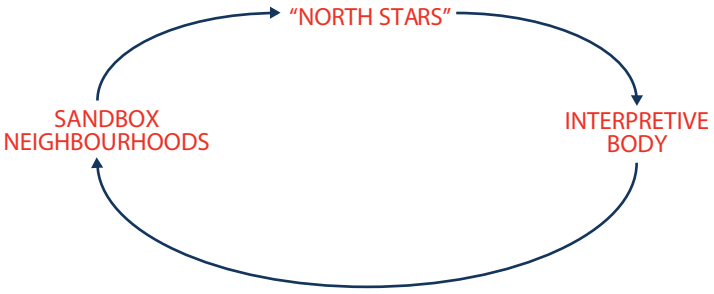

+ Develop "North Stars", or guiding values of AI, to inform policy and ensure it serves public interest goals. These would not have the force of law, but would instead be a normative framework[2].
+ Government should either create an independent body, or task existing bodies (such as the Residential Tenancy Dispute Resolution Service) that would be responsible for interpreting North Star values and assessing AI applications and their impacts within the housing domain.
+ Creation of sandbox neighbourhoods that allow housing developers and residents to experiment with different variants on regulation (for example, on data capture) to assess benefits and drawbacks that can inform revisions to North Star values.

DOMAIN: JUSTICE

Participants within this group were concerned with prescribing limits to what kind of decisions legal AI assistants or applications are able to make, asking "where do we draw the line?" This included concerns over the kind of information available to these systems, specifically whether these systems would use social media data to make legal decisions. Participants were concerned with how such systems would be audited to mitigate bias and unfair decisions.

+ Government should prescribe limits to manage how, and to what extent, intelligent legal assistants or applications influence or make legal decisions.
+ Policymakers should address the potential impacts this kind of technology may have on legal professionals, and make re-skilling options available for those who could be displaced from their positions.
+ Government should develop policies that would encourage competition within this space, minimizing the existence of monopolies and the amount of money spent in foreign software, products, and services.

DOMAIN: EDUCATION

Participants in this group recognized the potential difficulty of enforcing a ban on technology within schools. They therefore maintained that educational technology providers should gain consent from users before tracking students' learning abilities in the interest of student privacy.

Participants within this group also acknowledged that intelligent educational products may become an issue of global competitiveness if other countries allowed or promoted the use of these within their own educational institutions. This factor would need to be accounted for when developing policy around AI in educational settings. Moreover, they think it is critical for policymakers to find the right balance between ensuring these systems are transparent and open source, while managing the privacy of end-users. They therefore recommended:

+ Educational technology should be opt-in.
+ Educators implement policies to control the use of digital technologies within schools, permitting the use of these technologies to specified time periods in hopes of mitigating the associated risks.
+ Government implements mandatory maintenance and open source policies that would ensure regular auditing and provide greater transparency.

2 While not explicitly mentioned in the discussion, it is possible that these "North Star" principles could draw on existing AI values frameworks, such as the Montreal Declaration.

DOMAIN: HEALTH

Participants within this group were concerned with the lack of research funding that has been dedicated thus far to the social implications of AI within Canada. This group also questioned the ownership and control of patient data throughout the training and deployment lifecycle of an AI application, and believed government should implement safeguards to ensure citizens have control over their health data. Participants also recognized the issues the federal government would face when developing a national policy related to healthcare, and agreed that provincial governments would be the primary vehicles for developing policy addressing the use of AI in the health sector.

+ Increase funding for research into the social impacts of AI.
+ The government should adopt a similar framework to the General Data Protection Regulation (GDPR) to help enforce data ownership and privacy.

GENERAL REMARKS

Like Toronto participants, participant feedback from Edmonton indicated the benefits of being able to connect and collaborate with policymakers with an interest in AI in their own city. In fact, most participants did not know each other before attending the lab, which made for great networking. Most significantly, participants highlighted an appetite for more events like this, which explore the social and political impacts of AI within Canada. Participants also appreciated the ability to provide policy recommendations at the end of the day. However, participants also noted the difficulty in developing policy recommendations that relate to both the current and potential future state. Many recognized that AI policy should be specific to the technological capabilities and its application, and therefore preferred to address one case study in particular. When asked what their key take away from the day was, most participants noted the need for policymakers to keep apace with the rapid development and implementation of AI in society, and ensure their policies are agile enough adapt to technological change over the long-term.

NEXT STEPS

CIFAR and BII+E will be integrating participant feedback to inform the design of our third AI Futures Policy Lab taking place in Vancouver on October 18, 2018. If you would like to participate in one of our upcoming labs in Vancouver, Ottawa or Montreal please contact Gaga Boskovic.

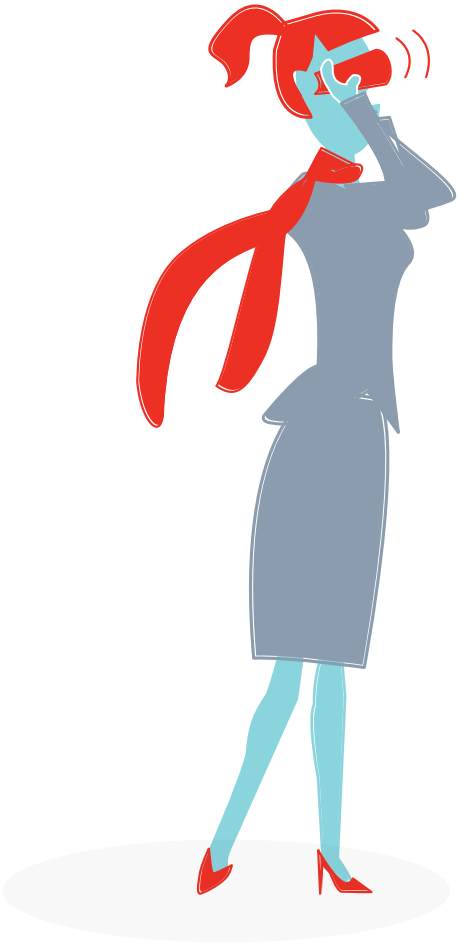

APPENDICES

APPENDIX A: AGENDA

Time	Total Time	Activity
8:30am	**30 min**	**Light Breakfast + Networking**
9:00am	**25 min**	**Opening Remarks** + Brent Barron, Director of Public Policy, CIFAR + Heather Russek, Director, Policy Innovation Platform, The Brookfield Institute for Innovation + Entrepreneurship + Kim TallBear, Canadian Research Chair, Indigenous Peoples, Technoscience, and Environment
9:25am	**20 min**	**Introductions + Warm Up Excrcise** In small groups, participants will play The AI Thing From the Future, an imagination game that encourages players to generate the most interesting, funny, or thought-provoking ideas for artifacts from the future.
9:45am	**60 min**	**AI 101** + Cam Linke, Interim Executive Director, Alberta Machine Intelligence Institute (AMII)
10:45am	**15 min**	**Break**
11:00am	**60 min**	**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.
12:00pm	**45 min**	**Lunch**
12:45pm	**75 min**	**Examining AI in 2028** In small groups, participants will explore the social, political, economic, and ethical dimensions of a future AI scenario.
2:00pm	**15 min**	**Break**
2:15pm	**60 min**	**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	**45 min**	**Presentations & Closing** Each group will have the opportunity to present their policy recommendation to the larger room.
4:00pm	**15 min**	**Social & Networking**

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 though 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: *SMART DEVELOPMENT*

A housing developer has included smart home technology driven by artificial intelligence in a large condo complex under development. Built-in voice, facial and gait recognition, as well as embedded sensors, provide security, automated services, predictive maintenance, and enforcement of condo by-laws.

Residents are able to opt-out of the installation or data collection by this private company, providing they are willing to pay an additional fee representing the market value of their data streams and compensation to their neighbours for the reduced accuracy of predictions.

The municipal government proposes a public-private partnership with the housing developer to bring the same model to a new affordable housing development. In exchange for anonymized data rights of residents, the developer will build high quality, affordable housing for half the cost of a typical development – a welcome prospect for cash-strapped municipalities. Residents will still have the option to pay to opt-out of data collection, but the fee is effectively insurmountable for residents that qualify for affordable housing.

APPENDIX D: *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 E: *SUMMARY JUDGEMENTS*

Based in part on the success and accuracy of legal summary and prediction technologies, a small but increasing number of administrative and quasi-judicial bodies are now using myJudgement, an AI application that makes initial judicial decisions by scanning uploaded legal documents, and finding similarities across a database of past cases and related judgements using natural language processing (NLP). These automated decisions are limited to relatively small matters, such as fines or ticket disputes, and do not play a part in the criminal justice system. In all cases, decisions made by myJudgement can be appealed in front of to a human judge.

However, some access to justice advocates have raised the possibility that myJudgement could serve a barrier to those who are unfamiliar with navigating bureaucracy and tribunals and who don't "optimize" their written and verbal statements in a way that is interpretable to myJudgement's software. Others are concerned that automating low-level decisions is creating a rationale for reductions in resources across the justice sector. Yet, supporters of myJudgement argue that both of these issues problems existed prior to the introduction of automation to law.

APPENDIX F: *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.

Nestor's software can also integrate with students' social network profiles and calendars to suggest study times and foster more effective study habits. For example, if a student has a tendency to watch YouTube videos at 11:00am on Sundays, Nestor can suggest that as a time for a study session instead.

Nestor encrypts, anonymizes and stores analysis data, but does not currently keep video footage or sell it to advertisers.

APPENDIX G: *ADAPTIVE EDUCATION*

An increasing number of services and devices, from websites to wearables to toys, seek to monitor and optimize children's learning abilities. These range widely in terms of proven efficacy, standards, cost, accessibility, and data portability. In response to the growing market of smart educational devices, school boards have taken different approaches to whether – and how – to integrate these technologies into the classroom. Some private and charter schools have piloted the integration of smart toys to aid in student learning evaluation and classroom-tracking, in which a software tracks each student's progress in real-time, to building profiles of strengths, weaknesses, as well as identify at-risk students.

However, public school boards have banned the use of these products in the classroom on the grounds that they may interfere with standardized evaluation criteria and can affect a teacher's ability to manage their students.

APPENDIX H: *INNEREYE*

Project InnerEye, a research initiative led by Microsoft, applies state-of-the-art computer vision and machine learning algorithms to automatically analyze three-dimensional medical CT (Computer Tomography) and MR (Magnetic Resonance) images to identify tumours and organs at risk.

The current processes of marking up radiology images is time consuming and expensive, with images often only marked up once before radiotherapy begins, and once again at the end of the treatment cycle. InnerEye serves to enhance the workflow of healthcare professionals, such as radiologists, surgeons, and medical physicists by analyzing images pixel-by-pixel to identify the exact position and size of the tumour, as well as the healthy organs that surround it. This enables healthcare professionals to more effectively plan a patient's radiotherapy strategy or surgery navigation.

By making this process more effective and cost efficient, InnerEye patients can potentially receive "adaptive radiotherapy", with scanning, image markup, and therapy planning being done after every treatment session. In doing so, InnerEye can help identify which type of treatment works best by monitoring changes in tumour size.

InnerEye has been trained on scores of images from past patients that have been marked up by experienced health professionals, meaning its system should perform as well as a leading expert every time. Nevertheless, doctors retain full control of InnerEye's system, and can make adjustments to the software at any time until they are completely satisfied with the results they receive.

InnerEye is currently being used by the UK's National Health Service for prostate cancer diagnosis and treatment, but could potentially benefit any health processes that use 3D imaging.

APPENDIX I: *AI-DRIVEN HEALTH MONITORING*

An increasing number of services and devices, from websites to wearables to toys, seek to monitor and optimize children's learning abilities. These range widely in terms of proven efficacy, standards, cost, accessibility, and data portability. In response to the growing market of smart educational devices, school boards have taken different approaches to whether – and how – to integrate these technologies into the classroom. Some private and charter schools have piloted the integration of smart toys to aid in student learning evaluation and classroom-tracking, in which a software tracks each student's progress in real-time, to building profiles of strengths, weaknesses, as well as identify at-risk students.

However, public school boards have banned the use of these products in the classroom on the grounds that they may interfere with standardized evaluation criteria and can affect a teacher's ability to manage their students.

APPENDIX J: *THE AI THING FROM THE FUTURE KEY*

ARC outlines the type of future that the "thing" comes from, and how far away it is from today. There are four types of Arc cards, each an umbrella for countless possible scenarios:

1. *Growth:* a future in which "progress" has continued
2. *Collapse:* a future in which society as we know it has come apart
3. *Discipline:* a future in which order is deliberately coordinated or imposed
4. *Transformation:* a future in which a profound historical evolution has occurred

TERRAIN is the thematic context or location where this object could be found in that future. *OBJECT* is the focus of your imagination - a specific cultural artifact that reveals something about how this future is different from today. *MOOD* suggests how it might feel to experience this thing from the future. *AI* indicates the technological capability or application that needs to be integrated in the artifact you create.

As an example, imagine you are presented with the five cards below:

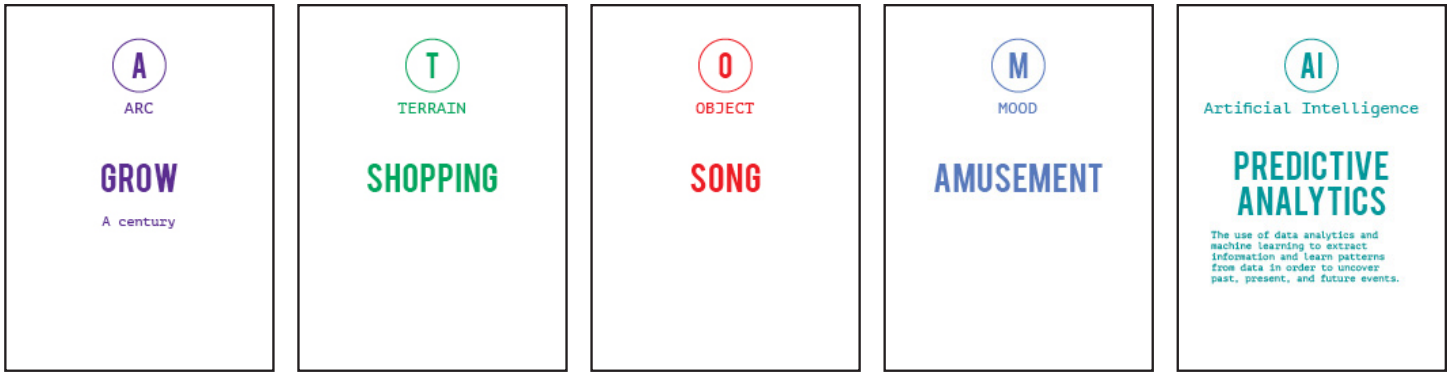

These cards point towards a future in which progress has continued, in the domain of shopping, with the focus being a song, accompanied by a feeling of amusement, and the use of predictive analytics. In imagining a thing associated with the prompts on these cards, you may think that a century from now, there will be fitting rooms that predict which songs you like to hear while you are shopping. This will help elevate the experience by leaving you with the feeling of amusement.

THE AI THING FROM THE FUTURE

brookfield institute for innovation + entrepreneurship

CIFAR

1) YOUR CARDS

ARC	TERRAIN	OBJECT	MOOD	AI
_____				
_____				

2) DESCRIPTION

_____

_____

_____

_____

_____

_____

_____

_____

_____

3) SKETCH

Adapted from Situation Lab (Stuart Candy and Jeff Watson)

APPENDIX L: *CANVAS 1*

Case Study.

Canvas #1: 2018

brookfield institute for innovation + entrepreneurship

CIFAR

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?

	Local	*National*	*Global*
Social			
Technological			
Environmental			
Economic			
Political			
Values			

Step 3 What existing policies and programs are affected?

APPENDIX M: *CANVAS 2*

Case Study.

Canvas #2: Envisioning 2028

brookfield institute for innovation + entrepreneurship

CIFAR

Step 1 What could have happened differently in this future?

Step 2 What are the potential impacts of this scenario?

	Local	*National*	*Global*
Social			
Technological			
Environmental			
Economic			
Political			
Values			

Step 3 How are different groups experiencing both positive and negative effects?

? Stakeholders	+ Positive	- Negative

APPENDIX N: *CANVAS 3*

Case Study.

Canvas #3: Responding Today

brookfield institute for innovation + entrepreneurship

CIFAR

Step 1	Step 2	Step 3
What are the most important positive and negative implications of this technology today?	What are the most important and negative implications of the scenario in 2028?	What stakeholder groups experience the strongest gains and the most significant losses.
+	+	+
-	-	-

Step 4 What policy responses today could help to improve outcomes today and prepare for potential futures?

Option

Benefits

Concerns

Step 5 What are your final policy recommendations?

1 2 3