

SHAPING AI

Attached is a draft version of the forthcoming report, A Northern Spark? AI Governance in Canada (2011 to 2021). The edited report compiles research conducted as part of the Shaping AI Canadian policy research activities that ran from 2021 to 2023. Research from this report was coordinated by Dr. Sophie Toupin and Dr. Fenwick McKelvey. The report is being submitted in draft to ensure the research can be consulted as part of the public record.

The report is part of a multi-national and multi-disciplinary social research project that examines the global trajectories of public discourse on artificial intelligence (AI) in four countries (Germany, UK, Canada, and France) over a ten-year period 2012-2021. Funded by the European Open Research Area initiative for a period of three years (February 2021 - February 2024), Shaping AI brings together leading research teams from each of the four countries under scrutiny.

Our research finds that:

1. Canadian AI governance has focused on economic and industrial policy.
2. National symbolic investment in AI – the promotion of AI’s Canadian-ness – impedes critical discussion about the technology and its risks.
3. AI governance is uncoordinated, lacking clear mandates for consultation and effective mechanisms of feedback that impedes good governance and public participation.
4. AI policy is marked by notable silences including Indigenous rights and data sovereignty, the creative and cultural sectors, and the environment.
5. The Government of Canada is a key site of AI development and deployment that remains understudied in the current legislation.

The views of each chapter are of the authors themselves. Collectively, the report raises a need to improve AI consultation and public participation in AI governance and seriously the assumptions and efforts that led to the development of the Artificial Intelligence and Data Act as well as its inclusion in the larger implementation of C-27. The passage of this act raises concerns not just about AI but about digital policy development and good governance.

Sincerely,

Fenwick McKelvey

SHAPING AI

A Northern Spark?

AI Governance in Canada (2011–2022)

About the Report

This report compiles research conducted as part of the Shaping AI Canadian policy research activities that ran from 2021 to 2023. Research from this report was coordinated by Dr. Sophie Toupin and Dr. Fenwick McKelvey. Thanks to Blair Attard-Frost and Aaron Tucker for providing peer review of the report.

This edited collection draws on research supported by the Social Sciences and Humanities Research Council.



Social Sciences and Humanities
Research Council of Canada

Conseil de recherches en
sciences humaines du Canada

Canada

Research Team

Shaping AI is a multinational and multidisciplinary social research project that examines the global trajectories of public discourse on artificial intelligence (AI) in four countries (Germany, UK, Canada, and France) over the ten-year period 2012 to 2021.

The Canadian research team is led by:

Fenwick McKelvey is Associate Professor in Information and Communication Technology Policy at Concordia University. He is the author of *Internet Daemons* (University of Minnesota Press, 2018), winner of the Canadian Communication Association's 2019 Gertrude J. Robinson Book Prize.

Jonathan Roberge is Full Professor at the Institut National de la Recherche Scientifique (INRS) in Montréal. He funded the Nenic Lab as part of the Canada Research Chair in Digital Culture, which he has held since 2012. His most recent edited volumes include *Algorithmic Culture* (Routledge, 2016) and *The Cultural Life of Machine Learning* (Palgrave, 2020).

Sophie Toupin is assistant professor in the Department of Information and Communication at Université Laval in Quebec City. From 2021 to 2023, she was a postdoctoral researcher at Concordia University (Horizon program) where she examined critical perspectives on artificial intelligence.

Guillaume Dandurand is a postdoctoral fellow at the Institut National de la Recherche Scientifique, in Montréal. He is coeditor of *Les économies de la promesse* (Presses de l'Université de Montréal, 2022). His doctoral research was awarded York University's Dissertation Prize (2019) and was included on the shortlist for the International Convention of Asia Scholars Social Sciences Dissertation competition (2021).

The team includes (in alphabetical order):

Marek Blottiere is a research assistant for the Shaping AI research project. As part of his FRQSC-funded master's degree in cultural studies at the Institut National de la Recherche

Scientifique, Blottiere investigates Montréal's media ecosystem and how it shapes the cultural politics of AI.

Nicolas Chartier-Edwards est étudiant au Doctorat Sur Mesure à l'Institut national de la recherche scientifique (INRS) et détient une maîtrise en sociologie de l'Université Laval. Ses travaux portent principalement sur l'intégration des technologies d'intelligence artificielle dans la production du social et plus précisément, dans le cadre de sa thèse, sur la transformation de la gouvernance canadienne par le déploiement de ces mêmes technologies dans les fonctions administratives des différents paliers étatiques, sois municipaux, provinciaux et fédéraux.

Nick Gertler recently earned his master's degree in media studies from Concordia University. Funded by scholarships from SSHRC and FRQSC, his thesis focuses on algorithmic impact assessments and algorithmic governance in Canada.

Etienne Grenier is an artist and researcher working in the field of digital cultures. Currently a PhD candidate at the Institut National de la Recherche Scientifique in Montréal, he studies the impacts of datafication on cultural production. Maintaining an active creative practice in digital arts, his installations and performances have been presented in leading institutions and major festivals in Europe and the Americas.

Robert Hunt is a PhD candidate in communication at Concordia University, Montréal. His SSHRC-funded doctoral research investigates the implementation of artificial intelligence in human resources and human capital management. His work has been published in *Social Media + Society*, the *Journal of Information Policy*, and the edited collection *Affective Politics of Digital Media* (Routledge, 2021).

Maurice Jones is a curator and critical AI researcher based in Tiohtià:ke/Montréal and Tokyo. He is a PhD candidate at Concordia University in Montréal, associated researcher at the Humboldt Institute for Internet and Society in Berlin, and the Artistic Director of the MUTEK.JP festival of digital creativity in Tokyo. His research interest lies in critical investigations of technology policy, crosscultural perceptions of artificial intelligence, and curatorial practice as research-creation.

Robert Marinov is a PhD student in communication at Concordia University with an MA in political science from the University of Ottawa, where his thesis won the Commission on Graduate Studies in Humanities Thesis Prize. His research focuses on the use of emerging Digital Twin platforms for governance and sustainability purposes, as well as their intersections with artificial intelligence. His work has been published in journals including the *Canadian Journal of Political Science*, *Communication Review*, *Politics & Policy*, and *Critical Studies in Media Communication*.

Meaghan Wester holds a SSHRC-funded MA in media studies from Concordia University; her thesis analyzed AI governance through public procurement. Wester is the winner of the 2022 CRTC Prize for Excellence in Policy Research. She uses the lenses of feminist infrastructure studies and feminist STS across her different research endeavours.

Table of Contents

Introduction by Fenwick McKelvey, Sophie Toupin, and Maurice Jones	6
Chapter 2: Canada’s Algorithmic Impact Assessment by Nick Gertler	32
Chapter 3: AI Source List: The Ethics of Procurement, the Display of Responsibility, and the Production of Nonknowledge by Meaghan Wester	43
Chapter 4: Consultations on Facial Recognition Technologies by Maurice Jones and Fenwick McKelvey	57
Chapter 5: AI and the Canadian Institute for Advanced Research by Robert Hunt and Théo Lepage-Richer (with research assistance from Nick Gertler)	79
Chapter 6: Chinook’s Nested Black Boxes by Nicolas Chartier-Edwards, Marek Blottière, and Sophie Toupin (<i>en français</i>)	103
Appendix by Sophie Toupin and Fenwick McKelvey	116

Introduction by Fenwick McKelvey, Sophie Toupin, and Maurice Jones

The joke goes like this: Canada is three oil companies in a trench coat. Or three telecommunications companies. Unchecked corporate power is not a new issue in Canada, and the country does not have a successful record solving it. The artificial intelligence industry has a similar problem, with a few large firms dominating the market and steering its future. These are global challenges. Our report sets out to explain how Canada has responded to the challenges and opportunities of AI. What we find is that the joke can be flipped: AI is Canada in a trench coat. Rather than finding exclusively novel governance issues arising with AI, what we instead see are the enduring problems of Canada, or at least the federal state, reappearing in a new guise. After three years of researching AI governance in Canada, our story is less one about technology than one of the nation. Understanding AI governance in Canada is a study in federalism, history, and dependency. It is a story of national anxieties and local ambitions. It is not a story about technological creation but of myth-making—an investment in building a better workforce, a better government, a better Canada.

As with the nation, coherence is elusive. The challenge with trying to tell a story about AI governance is to communicate the multitude of parts—especially when the prevailing narrative is so simple. The usual story begins with the newly elected Prime Minister Justin Trudeau, son of former Prime Minister Pierre Elliot Trudeau. Both had an infatuation with technologies and statecraft (Lepage-Richer & McKelvey, 2022). As the “Instagram Prime Minister,” Justin Trudeau’s early press conferences portrayed the leader as an enthusiastic geek, apt to talk about quantum computing and artificial intelligence. AI was a classic Canadian underdog success story. Machine learning was an approach that ran counter to the mainstream investment in expert systems, yet Canada invested in machine learning nonetheless, training and teaching the leaders of the AI revolution: Geoffrey Hinton, Yann LeCun, and Yoshua Bengio. Eventual winners of the prestigious Turing Prize, their success was made possible by Canadian government investment (Metz, 2019). When Trudeau announced the launch of the first round of its Pan-Canadian AI Strategy, he did so at the elite University of Toronto, touting both his geek credentials and this homegrown Canadian victory. There can be no

doubt that Canadians have contributed massively to the research behind today's AI summer (CIFAR, 2017c), but how these accomplishments have been managed and how AI affects society and Canada has been less clear.

How can we tell a more complicated story of how Canada shaped AI? To start, we offer something like what Anna Tsing (2015) calls fast stories, which, instead of trying to add things up, find meaning in the differences among each story. These stories help to situate our findings and our report.

2017: Canada announces its pan-Canadian AI strategy. Canada becomes one of the first countries in the world to commit to a national plan for AI development, initially consisting of a \$125 million federal investment concentrated in three cities: Edmonton, Montréal, and Toronto. The government tasks the Canadian Institute for Advanced Research (CIFAR) with managing the strategy, building on the organization's long history in AI research. Founded in 1982 as a nonprofit research institute designed to foster interdisciplinary and inter-institutional networks among researchers, CIFAR had been involved in AI research since the institute's inception. But its prominence in the field increased dramatically in the decade preceding the strategy, due to program leader Geoffrey Hinton's success in using neural networks to dramatically improve image recognition and in attracting the interest of major tech companies like Google and Facebook.

2018: Concern about facial recognition technologies (FRTs) started at the mall. "Facial recognition at Chinook," read a post in the Reddit community /r/Calgary. "Chinook" meaning the Chinook Center, a mall in downtown Calgary owned by real estate firm Cadillac Fairview. The next month the Privacy Commissioner of Alberta (OIPCA) received a formal complaint triggering an investigation into the technology's use in the province. Government inquiry into FRTs—referred to as anonymous video analytics at the time—began with little public notice. Facial recognition slowly grew into a national controversy after reporter Kashmir Hill broke the Clearview AI story on 18 January 2020 in the *New York Times*. Hill's story reported that the American firm had trained its facial recognition system using billions of images uploaded onto social media; he

also mentioned that unspecified branches of Canadian law enforcement had used the tool. A few days later, national broadcaster CBC’s flagship radio program the *Current* aired an interview with Hill that failed to detail Canadian experiments with FRTs (“The End of Anonymity?,” 2020). Even with warnings from the Canadian Civil Liberties Association’s Brenda McPhail that “we’ve got law enforcement agencies using [FRT] without confirming that the use of the tool is compliant with Canadian law. That’s a big problem,” the program concluded with the assertion that no police forces used the technology. A few days later, possible uses in Canada did not receive a mention. On 1 February 2020, a CBC story on Clearview by Ramona Pringle (2020) neglected to mention a single Canadian case. That, of course, was wrong.

2020: After several years of development, Canada’s algorithmic impact assessment (AIA) comes into force as part of a policy directive on automated decision-making. The AIA “helps designers understand and manage the impacts of their AI solutions from an ethical perspective” (Government of Canada, 2024). The tool, functionally a platformized questionnaire used to assess risk, becomes part of the Government of Canada’s strategy for responsible AI deployment along with new rules for the procurement of technology. The government tells the AIA’s story differently than we do; its website presents a timeline of the project that moves smoothly from design to execution—an uncomplicated story of governmental foresight and intervention. Our report has no such line because we found a much messier history. Missing from the government’s account is the remarkable proliferation of alternative AIAs in the late 2010s, the result of something of a public-private cottage industry that produced a variety of competing surveys and tools. This messier story lives on in the official, open-source AIA—developed by the government but maintained as a GitHub repository—and in the corporate consensus on “responsible AI,” a discourse that began years before the government started work on regulating AI.

2021: French-language media outlets report that a new immigration AI system—named Chinook—refuses a greater number of francophone international students. This other Chinook generates a small scandal and calls to learn more about the role of technology in Canada’s immigration system. The story becomes largely one about the problems of artificial intelligence. Activists closer to the system tried to leverage interest in AI:

If, for the umpteenth time, I repeat, “Oh, the immigration process is slowed down because Quebec and Ottawa do not agree on the thresholds,” after fifteen times, the journalist will no longer write this article. So we went looking for a “buzz word” that nobody knew about. So we went to ask the experts for their opinion and there we opened a whole new chapter [through the notion of artificial intelligence]. . . . We knew, I knew, what AI was [and wasn’t]. But just the key word, even if it wasn’t really [AI], it allowed us to make some mileage out of it. (X, personal communication, September 19, 2022)

Immigration institutions are rife with systemic racism and the process of applying for immigration is dehumanizing, but AI gets more media attention. Immigration Canada continues to invest in artificial intelligence solutions despite a growing concerns that internal biases might become encoded in its systems (Keung, 2023).

2022: Canada launches a new Artificial Intelligence and Data Act, catching most people by surprise. The government tabled AIDA as an addendum to its omnibus consumer privacy act. Logically this report should include the story of AIDA’s development, but we were as surprised as anyone. AIDA is not an endpoint, or a culmination, but another moment in the fragmented history of Canada's approach to AI. Thus, our period of study from 2012 to 2022 ended in an awkward, but in many ways appropriate, spot.

What to make of these fast stories about AI in/and Canada? AI in some ways becomes less important as a technology than a particular site of intervention for Canada’s persistent nationalist insecurities to play out.

Our project fits with a growing turn toward the study of AI governance and, more specifically, Canada’s federal AI strategy. Many of the international reviews of AI strategies have consisted of global scans, reviewing the stated or discursive framings of AI as well as specific policy measures to promote AI, usually under the umbrella of industrial policy (Bareis & Katzenbach, 2021a; Cath et al., 2018; Hansen, 2021; Köstler & Ossewaarde, 2022; Kuziemski & Misuraca, 2020; Margetts, 2022; Paltieli, 2021; Radu, 2021; Sun et al., 2020; Taeihagh, 2021; Ulnicane et al., 2021). We deviate in two ways. First, by taking a national view, we include subnational or provincial differences (cf. Jobin et al., 2019a). A global approach, while helpful for comparison, misses local dynamics, such as how parallel provincial and federal AI strategies have synergized to turn Montréal into an AI hub. Our research into facial recognition consultations in Canada, as another example, demonstrates how consultations at the local, provincial, and federal level can compound. Secondly, our approach differs through its attention to AI on the ground. Drawing on lessons from Tess Lea’s (2020a) research on policy haunting, we choose to extend our literature review into tangible cases.

We enjoy being part of a growing national review of AI policy. Ana Brandusescu published a landmark study of the funding priorities for public AI investments, finding that “public investments in AI technologies primarily benefit the private sector” (2021, p. 7). Our research complements Brandusescu’s review by similarly considering the Canadian Institute for Advanced Research’s influence, but instead analyzing how the nonprofit has become the hub of its own network. More recently, Blair Attard-Frost, Ana Brandusescu, and Kelly Lyons (2023) conducted a scan of 84 government initiatives, including projects in Alberta, Ontario, and Quebec. Our research supports their study’s findings that Canada’s overall AI policy is uncoordinated and lacks consultation innovation. Finally, as noted previously, our research concluded before the introduction of the Artificial Intelligence and Data Act (AIDA), but there is now growing discussion of AI governance in Canada, post-AIDA. We offer a longer story of what happened before AIDA.

Building on grounded research (see Appendix), we first focused on document and policy analysis to understand the relevant languages and discourses. We find:

1. Canadian AI governance has focused on economic and industrial policy.

2. AI is intertwined with the Canadian and Quebec governments' self-promotion and branding, making it a nationalistic technology.
3. AI governance is uncoordinated, lacking clear mandates for consultation and effective mechanisms of feedback.

From these general trends, our research turned toward:

1. How artificial intelligence is being used as a new regulatory and service delivery tool in governments;
2. How Canada's industrial AI policy has emerged; and,
3. How AI consultations inform AI policy development.

These questions become the basis for the report's subsequent chapters; this introduction elaborates our key points. Our choices, subsequently, map onto Regine Paul's guidance for critical policy research on AI. Paul argues that a critical policy agenda for AI studies has a

triple focus on (1) the interpretivist reconstruction of values and meanings in AIT-policy-interactions, (2) the exploration of concrete practices of sociotechnical co-production, and (3) critical discussion of how AIT uses in policy interact with (global) inequality and wider power relations—seeks to set a scholarly counterpoint to techno solutionism. (R. Paul, 2022, p. 505)

Our cases build on this agenda.

Canada as an AI Industrial Superpower?

Canada has positioned itself as an international leader in AI, but, in reality, activities have been concentrated in a few major cities. The emergence of Montréal and Toronto as AI hubs started around the mid-2010s but is rooted in a longer history that goes back to the 1980s and early 1990s. Geoffrey Hinton and Yoshua Bengio arrived, respectively, at the University of Toronto in 1987 and the Université de Montréal in 1993 and were to become known as the “AI godfathers” (Senneville-Robert, 2021). CIFAR funded their earlier work. In 2012, Hinton won the ImageNet prize, which catapulted the connectionist approach to AI to notoriety. Hinton, Yann LeCun, and Bengio then received the Turing Award in 2018 for pursuing research on deep learning. Hinton and LeCun accepted positions in the business world, while Bengio

created the Montréal Institute of Learning Algorithms (Mila) and, in 2016, the firm Element AI.

The narrative of Canada as an AI supercluster is much more complicated on the ground. The international reputation of Canada's AI researchers has found itself immersed in local, provincial, and federal efforts to accelerate AI as an industrial opportunity and to position itself in a shifting global landscape.

A federal strategy

Canada developed a robust network of coordinated AI institutions during our period of study (Senneville-Robert, 2021). In 2015, when the current Liberal government was first elected, discussions were being held to keep Canadian AI talent in Canada (what we call AI nationalism). In Toronto, the cofounders of Layer6 convinced Ontario's provincial Liberal government to fund a new research institute on AI, the Vector Institute, with \$50 million. In 2016, the federal government allocated \$93.6 million to the Institute for Data Discovery (IVADO), which further accelerated the development of Montréal and Toronto as AI hubs.

These efforts culminated in Canada launching the first national AI strategy worldwide. Many countries would follow suit. In 2017, the Pan-Canadian Artificial Intelligence Strategy (PCAIS) was set up by the federal government with an investment of \$125 million. CIFAR supervised the strategy and helped position the Edmonton-based Alberta Machine Intelligence Institute (Amii), Montréal-based Mila, and Toronto-based Vector Institute as the three pillars of the Canadian AI pipeline (Senneville-Robert, 2021). Two years after its launch, the government extended oversight of the PCAIS with the launch of its Advisory Council on Artificial Intelligence as well as a Commercialization Working Group. These initiatives sought to bolster Canada's strength in research and attract talent.

National developments paralleled Canada's more normative presence on the international stage. As part of its G8 (now G7) presidency, Canada launched the Global Partnership on AI in 2020. This multi-stakeholder initiative coordinated global discussion on AI's social impacts, especially Responsible AI, and published a number of reports and recommendations. Other international efforts tended toward bilateral or multilateral agreements, such as the 2019

Canada-UK Artificial Intelligence Initiative research funding agreement or more adjacent statements such as the Christchurch Call to Action Summit or UNESCO's guiding principles on the diversity of content online (Canadian Heritage, 2021; Public Safety Canada, 2022). Notably, we struggle to find a connection between these initiatives and national or provincial action in Canada.

As Canada positioned itself as an international leader on AI strategy, doubts lingered about its independence. Critics of the PCAIS have raised concerns about its lax approach to intellectual property and further enmeshing of Canada's technology sector with large US firms. As the former chair and co-chief executive officer of the Canadian technology company, Research in Motion, Jim Balsillie explained to the Industry, Science and Technology Committee in 2020,

Canada is on the sidelines in the global competition for IP and data, contributing to their creation but not contesting their ownership and ensuing benefits. Consequently, we see the exfiltration of knowledge assets out of Canada on a regular basis, across borders with the stroke of a pen, currently without any national security or economic review. For example, foundational IP for AI that Canadian taxpayers have funded for two decades is transferred from the University of Toronto to Google. (Balsillie, 2020)

This warning came true when US firm ServiceNow acquired Element AI—arguably Canada's flagship AI start-up—in 2020 for cheap after its value collapsed (Silcoff, 2020). Despite these concerns, the government continued to support this initiative with \$125 million included in the 2021 budget over five years, from fiscal year 2021–2022 to 2025–2026.

Canada's AI strategy has been an industrial strategy first and foremost. In May 2017, the federal government and then-Minister of Innovation, Science and Economic Development Navdeep Bains announced the allocation of \$950 million to fund the Innovation Supercluster Initiative (Balingall, 2018; Doloreux & Frigon, 2022). This initiative was aimed at creating “more middle-class jobs and more opportunities for Canadian businesses to grow into globally successful brands.” The approach borrowed from conventional cluster industrial policy, launching five “superclusters” in 2018. Initial members included research centers, such as IVADO, universities, such as the Université de Laval, HEC Montréal, and the University of Toronto, and local AI players, such as Element AI and Mindbridge AI. Industry players

included firms in transportation (Air Canada and ViaRail), pharmaceuticals (Pharma Science), technology (Opentext, Intel, and Cisco), banking (Desjardins, BMO, and RBC), telecommunications (Bell), retail (SSENSE and Shopify), as well as venture capital firms, start-up incubators, and consulting firms.

These efforts have led to moments of policy debate outside the usual channels. Beginning in 2019, the Canadian Radio-Television and Telecommunications Commission (CRTC), Canada's media regulator, granted indefinite approval to Bell Canada to use an automated tool to block malicious calls on its network and any network transiting its networks (McKelvey et al., 2022). In doing so, the CRTC set important precedents for AI governance, principally by rejecting the need for public explanation of how the systems work, instead putting security first at the cost of transparency.

Despite the government's investment, public attitudes to AI are reserved. A 2019 study from Deloitte found that 65% of Canadians had "privacy concerns over how companies were using their data today" (Deloitte, 2020, p. 4). Globally, Canada has the second lowest (32%) support for the statement "Products and services using artificial intelligence have more benefits than drawbacks" according to a 2022 Ipsos poll for the World Economic Forum (Government of Canada, 2021a). A 2021 survey done for the AI Advisory Council found that 65% of Canadians believed that "AI has the potential to cause harm to society," with only 48.2% of Canadians believing that "People who develop AI do so in an ethical manner" and only 43.2% agreeing that "Computers can be programmed to make ethical decisions" (Ipsos, 2022). Of course, with AI literacy being low according to most studies, a considerable gap exists between awareness, trust, and Canada's vision for AI.

Responsible AI in the federal service

The Canadian government, for its part, has tried to lead efforts in responsible AI by example. Its internal initiatives center on a Directive on Automated Decision-Making and a Trusted Suppliers List for AI vendors. Adoption of both procedures has been sporadic, and the tools have major shortcomings as we discuss in the later chapters. The Trusted Suppliers' List merely fast-tracks applications and does not include a robust vetting process for responsible

development. The Directive on Automated Decision-Making, by contrast, expects departments to conduct an Algorithmic Impact Assessment before deploying a system, but the assessment gauges risk without clear repercussions for projects deemed high risk. These efforts have not developed into a coherent national strategy.

Federal efforts to integrate artificial intelligence began soon after the Pan-Canadian AI Strategy. A presentation from 2018, obtained through an Access to Information Request, listed ten federal ministries developing AI pilot projects.

Early successes

Department	Use Case
NRCAN	<ul style="list-style-type: none"> • Forest fire protection and flooding map • Optimizing industrial process and performance • Extraction of features from high-definition images
TC	<ul style="list-style-type: none"> • Risk-based oversight of air cargo
PHAC	<ul style="list-style-type: none"> • Early warning of worldwide public health threats (in collaboration with NRC)
HC	<ul style="list-style-type: none"> • Automation of repetitive rules-based tasks • Assessment of risk communication activities
CSA	<ul style="list-style-type: none"> • Next generation holographic console
CRC	<ul style="list-style-type: none"> • Management of spectrum • Engineered surfaced design
ESDC	<ul style="list-style-type: none"> • T4 reissuance automation, including handwriting recognition • NLP-based risk determination for internal audit
ECC	<ul style="list-style-type: none"> • The creation of land cover information from satellite data • Improved and automatic classification of sea ice • The characterization, mapping, and monitoring of ecosystems and habitats important to wildlife
NRC	<ul style="list-style-type: none"> • Tracking greenhouse gases in the port of Montreal • Teaching machines disappearing indigenous languages • Helping to manage power consumption in buildings
VAC	<ul style="list-style-type: none"> • Determine documents of corporate value for archiving

Figure 1. List of AI projects in government in 2018

Where AI has been reported to be used in immigration and at the Department of Defense, the table above demonstrates AI’s application across the government with a clear emphasis on risk, threat, or value analysis—all examples of what might be called predictive analytics. These

examples serve to emphasize how AI has long been a project to reform government services as well as an industrial strategy.

Social AI policy?

Efforts to develop social policy for AI developed more informally. In 2017, a forum on Socially Responsible Development of AI was held in Montréal. This forum launched an eighteen-month consultation that led to the 2018 Montréal Declaration for a Responsible Development of AI. The nonbinding statement was an influential attempt to cultivate AI ethics particularly in the Montréal ecosystem. Quebec further invested in Ethical AI. In 2018, the International Observatory on the Societal Impacts of AI and Digital Technology (OBVIA), a consortium based at Université de Laval, was funded by the Quebec government.

Provincial Efforts

Quebec and Ontario are reforming their data and privacy laws to better accommodate AI, while the federal government has promised to pass its own privacy reforms. The Quebec government has independently invested in its own AI strategy, making Montréal the target of combined provincial and federal support. While the *Grappe industrielle en IA du Québec* (Quebec industrial AI cluster) channels most of its energies toward the Université de Montréal, the institutional and political tone is mostly one of justification, where “responsible” becomes an unquestioned leitmotif (Senneville-Robert, 2021). The Ontario government has started an extensive consultation around its AI governance and data strategy (Government of Ontario, 2021). These efforts continue in parallel to federal strategies.

Key Components of AI Governance

We offer a sense of the highly complicated, multi-level processes claiming to participate in Canadian AI governance in Figure 2. Our diagram maps AI initiatives from international to local (vertically) and from highly concrete actions, such as policy and law, to abstract, normative, or ethical statements (horizontally). The General Data Protection Regulation (GDPR), for example, is a concrete international law significant to Canadian actors worried about maintaining compliance. At the opposite end, we can identify the Montréal Declaration,

which was a local initiative to establish broad, normative goals around responsible development but without concrete enforcement or accountability. This diagram, to be fair, is an abstraction of the sites or moments of debate and discussion during our period of observation.

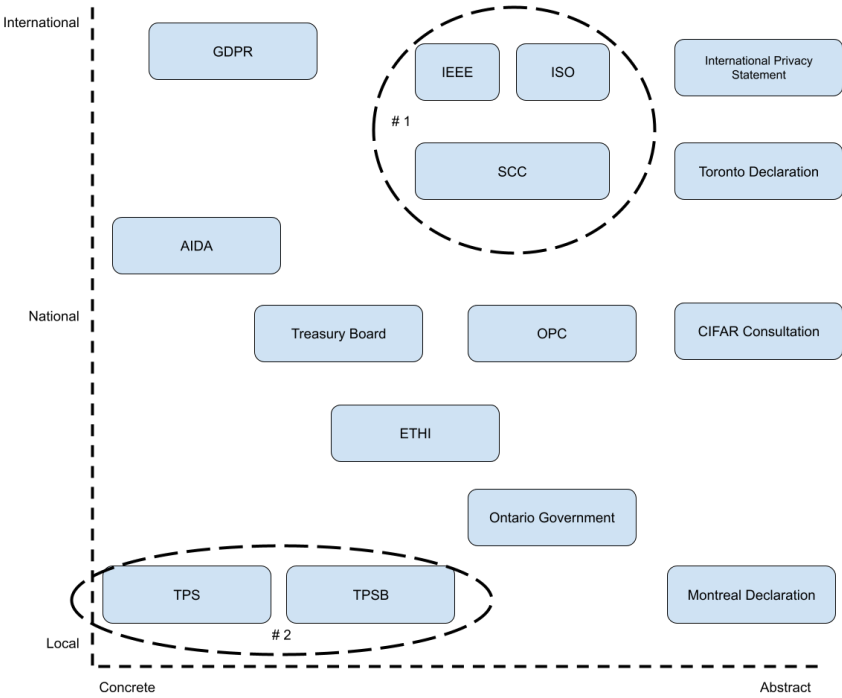


Figure 2. The governance assemblage in Canada

We can identify three key clusters of activity seen in Figure 2:

1. International Standards (cluster #1): The Institute of Electrical and Electronics Engineers (IEEE) and the International Standards Organization (ISO) are two key organizations behind the development of AI standards. The enforceability and acceptance of standards is an open debate, but the Standards Council of Canada (SCC) is coordinating Canada’s participation in ISO supported by the Pan-Canadian AI Strategy. Standards did not actively appear in our documentation, and they raise key questions about governance as participation in standard making is not usually public and participants are vetted.

2. Toronto Police Service (cluster #2): The TPS, discussed in a later chapter, initiated a public consultation around its AI procurement policies. Given the extensive use of AI in Canadian police forces, cluster #2 highlights some of the missing parts of the map, where the process of AI governance goes unseen.
3. Declarations: Much of the AI governance discourse focuses on declarations or ethical statements. The two most prominent being the Montréal Declaration and the Toronto Declaration, the latter of which is much less cited but represents an important moment of collaboration on responsible AI in international civil society.

In sum, we find Canada's approach to AI's societal implications, to borrow a phrase from policy scholar Tess Lea, "vertiginous and incoherent" (2020a, p. 25). Well, not entirely—AI is clearly more coherent as an economic priority despite growing social concerns and low levels of trust and acceptability. We find more of a policy assemblage. Clarke et al. (2015) describe policy assemblages as "elaborated, developed, and deployed in specific times and places as 'solutions' to the defined problems of those times and spaces" (J. Clarke et al., 2015, p. 50). In other words, rather than trying to paint a complete picture of Canadian AI policy in its entirety, what we offer is a view of different moments as parts of the assemblage.

Narratives



Figure 3. A word cloud of key terms in narrative coding

Policy documents express desires and participate in larger sociotechnical imaginaries about AI (Bareis & Katzenbach, 2021b). Rather than revealing a single dominant narrative, we found a strong cluster of narratives describing AI as positive as we elaborate in the Appendix, benefiting the economy and society as well as offering greater efficiency. To a lesser extent, numerous codes demonstrate some of the drivers of AI including trust, risk, and transparency. Notably, the geopolitical threats of AI did not appear as frequently as a positive language of opportunity and Canada's global leadership. Competition also does not appear significantly in the documents, further marking a notable difference between Canada and global narratives of an AI arms race.

Narrative 1: Shaping a good technology

Responsible AI represents a specific Canadian framing of a rapidly growing corpus of declarations published by national governments, private corporations, and civil society actors loosely summarized under the umbrella term of AI ethics (Jobin et al., 2019c). Initiated by the Canadian government, “responsible AI” became the Canadian AI ethics brand deployed by a

wide range of private and public (mostly federal) actors, including joint industry-academia-led initiatives, such as the Montréal Declaration on Responsible AI, and alternative civil-society-led initiatives like the Toronto Declaration. In addition, while under Canadian chairmanship the G7 meeting established the Global Partnership on AI, which specifically referred to “responsible AI” as its guiding principle. As it became a driving force in federal and global framings, responsible AI was also adopted as a guiding principle by Canadian provinces. Most notably, Ontario framed their approach as “trustworthy AI.” The Quebec government originally opted for “ethical AI” but later shifted to responsible AI.

The government of Canada began to cultivate the term in October 2016 during the drafting of a whitepaper on the “Responsible Use of AI.” As a government website on the responsible use of AI now declares:

Artificial intelligence (AI) technologies offer promise for improving how the Government of Canada serves Canadians. As we explore the use of AI in government programs and services, we are ensuring it is governed by clear values, ethics, and laws. (Government of Canada, 2024)

Responsible AI calls for leveraging the potential benefits of AI in the case of improving government services, while controlling the risks of the technology. The government of Canada’s Directive on Automated Decision-Making is a key text in the promotion of responsible AI. The directive raises concerns surrounding questions of bias and other harms, while centering opportunities for increasing government efficiency and improving service delivery.

The introduction and application of AI is never questioned; the technology simply arrives with the promise of being responsible. This sense of inevitability around the technology and its capacity to transform society necessitates getting things right from the get-go.

Responsible AI enables AI to become a source of national pride and a rearticulation of technological nationalism. Past work on technological nationalism focused on the railroad and

television; AI has similarly become a site of national significance and part of the performance of nationhood. In contrast to those earlier technologies, however, AI is deeply symbolic and less infrastructural. This form of technological nationalism is premised on constructing AI as a benevolent technology that reflects Canada's status as a benevolent nation. The Vector Institute, for instance, frames itself as a credible Canadian institution whose "Canadian values" guide it in developing responsible and beneficial uses of AI. In this context, the alternative framings of AI among the provinces, such as Ontario (trustworthy) and Quebec (ethical), function within their own nation-building projects as a way to differentiate themselves from the Canadian federal government.

Canada's efforts to co-construct itself with AI are consistent with international trends. Bareis and Katzenbach (2021b) showcase how national AI strategies set rules, allocate resources, and manifest sociotechnical imaginaries around the importance of AI for the nation-state. What our research shows is that this nation-building aspect of AI as culture industry might have been initiated by the government, but it is upheld and spread by a variety of public and private actors who have also hopped aboard the responsible AI train.

Narrative 2: Urgency of pushing legislation through

AI governance is urgently needed. The widespread uptake of AI as a governance issue underlines a phenomenon that Smuha (2021) describes as:

the "race to AI" is also bringing forth a "race to AI regulation." Indeed, a new playground for global regulatory competition seems to emerge, which in the best-case scenario pushes governments—amidst uncertainty as to the technology's impact, the impact of regulatory intervention, and the cost of non-intervention—to find the most appropriate balance between protection and innovation. (2021, pp. 59–60)

The AI regulation race is not only directed towards the Canadian context but towards establishing global leadership vis-à-vis other countries and regions. Most telling are the frequent references in Canadian policy documents to international attempts at digital and AI

regulation. The European GDPR and emerging AI Act are frequently mentioned in Canadian policy documents, as are regulatory efforts in places, such as the UK, Japan, and France, that push their own brand of ethical and human-centered AI. These policy documents seldom refer to Canadian federal or provincial efforts, despite the uptake of AI governance by a variety of government branches.

The shift towards a race to AI regulation is further supported by frequent references to the urgency of regulating AI in the wake of the transformative impacts it will have on society. For instance, in calling for consultations on a Modern Copyright Framework for Artificial Intelligence and the Internet of Things, federal government department Innovation, Science and Economic Development Canada (ISED) underlined the urgency of updating existing copyright legislation, which is inadequate to address rapid technological advancements and the questions they pose for copyright law. ISED's call states:

Since the beginning, copyright law has been shaped by major technological developments, starting with the printing press that led to copyright law, and then by other new technologies, from photography, to radio broadcasting, photocopiers and home recording devices, to the Internet and digitisation. To keep pace with these technological disruptions, copyright law has had to respond and adapt.(Government of Canada, 2021b)

This urgency is framed as the will of the Canadian people. The Honourable David Lametti, Minister of Justice and Attorney General of Canada in the announcement of the consultation on harmful online content explained:

Canadians expect that their government will protect them from harmful content online. The consultation held over the summer was an important step toward creating a safer and more inclusive online environment for all Canadians. The report being released today will help us develop a new framework for addressing harmful content online,

while respecting freedom of expression in a free and democratic society.(Canadian Heritage, 2022)

Yet, as a US campaign to ban facial recognition spearheaded by a dozen technology, media, and civil liberties organizations points out, often the urgency to pass regulation is driven by the cynicism of industry players. As their website states:

Like nuclear or biological weapons, facial recognition poses a threat to human society and basic liberty that far outweighs any potential benefits. Silicon Valley lobbyists are disingenuously calling for light “regulation” of facial recognition so they can continue to profit by rapidly spreading this surveillance dragnet. They’re trying to avoid the real debate: whether technology this dangerous should even exist. Industry-friendly and government-friendly oversight will not fix the dangers inherent in law enforcement's use of facial recognition: we need an all-out ban.¹

That final point deserves emphasis: despite superficial claims towards the responsible development of AI, in the end it is AI's economic potential that presents the main concern for governments and, consequently, policy. While policy documents pay lip service to responsible AI, the domination of industry-academia-government in governance as well as the resulting concrete measures and abundant allocation of resources tell a different story. AI becomes both a technology and industry where Canadian nationalism can maintain imaginaries of being ethical, innovative, and sustainable, while at the same time focusing on hard industrial policy. Canada is not alone in wrapping industrial policy in a language of ethics. Similar framings emerged in comparative cases, such as the making of Germany's national AI strategy (Jones, 2022).

These narratives coincided with Canadian media coverage of AI since 2012. Media narratives of AI frequently present it as a valuable resource brimming with potential for the local

¹ From the Ban Facial Recognition Website: Cite: <https://www.banfacialrecognition.com/>

economy and society. In the Canadian media, there's often an alignment with the federal and provincial governments' endorsement of AI as a collection of technological innovations with inherent "potential" that must be harnessed. Over the years, a substantial amount of public funding—over \$600 million—has been injected into AI research. However, Canadian AI coverage remains closely tethered to the investments made by businesses and governments in the AI sector. Our research findings indicate that the majority of AI coverage is concentrated within business sections, typically lauding the imminent capabilities of AI. An analysis of our chosen topics reveals that more than 45% of news articles focused on AI's future applications and its political economy compared to those addressing ethics and social debates. Critical perspectives are infrequently featured in AI coverage, as we found in our prior research. The same computer scientists repeatedly interviewed for AI stories continue to champion the expansion and implementation of various AI systems (Dandurand et al., 2022). In essence, our extensive two-year study of traditional media's AI coverage suggests that the ethical, social, and political dimensions of AI may not be perceived as newsworthy in the same way its economic impacts are.

Silences

Our research tried to find what was missing as well as what was present in the policy documents. We were influenced by situational analysis (Clarke et al., 2015) in helping our team of researchers identify the silences in the Canadian AI ecosystem. This is one of the strengths of situational analysis compared to grounded theory; it is designed to help silences speak by analyzing absent positions in the maps created. Situational analysis is intentional in “working against epistemic violence that erases or silences minor voices” (J. Clarke et al., 2015, p. 22). Table 1 lists the silences that were identified during our research discussed further in the Appendix.

Indigenous peoples
Montréal Declaration
Canadian Charter of Rights and Freedoms
Anti-racism

Toronto Declaration
Affected communities
TPSB
Resource allocation/tangible actions
Constructed Public
Clearview AI
Civil Society
Lack of Clarity
Impact assessment
Funding
Arts & Culture
Ambiguous data
AI risks

Table 1. Key silences in policy documents

Issues related to Indigenous people topped our sense of what was missing from these documents. Our approach to policy silences draws on Des Freedman (Freedman, 2010), who emphasizes the need to study negative policy, or governance involving what’s missing. What we add, drawing on situational analysis, is an attempt to find not only instances when policies are not discussed but other kinds of silences. For example, the policy documents did not consider approaches related to Indigenous data sovereignty. We point out this gap not as experts in this issue but as allies who view the supposed disruptive potential of AI as grounds to reconsider status quo policy approaches.

A second gap appeared concerning human rights and AI governance. Notably, despite explicit calls to adopt a rights-based approach from the Office of the Privacy Commissioner and others (McKelvey, 2018), the Canadian Charter of Rights and Freedoms was another key silence. This absence is remarkable because AI governance could readily be approached through interpretation of the Charter, rather than taking a negative policy approach that considers ethical AI to be a tabula rasa. Aspiration-based declarations also appeared infrequently in the documents. While the Montréal Declaration might have been an important statement for the local community, it did not have the effect of guiding approaches to AI policy. Finally,

antiracism was glaringly absent, even in the face of clear evidence of AI's discriminatory effects, calling into question broader statements of responsible AI (eg. Benjamin, 2019).

Civil society, publics, and affected communities were all missing from much of the language in governance documents. The Toronto Declaration, a collaborative statement whose authors included Amnesty International and Access Now, was an important early call for AI to defend human rights. Even though the statement had twenty-four international signatory civil society organizations, policy documents discussed and legitimated the Montréal Declaration far more than the Toronto Declaration. We discuss this issue more broadly in subsequent chapters, but through our research we found a gap between AI promotion and those affected, even the perceived beneficiaries. Civil society and communities were active during our study, especially in public committees and consultations, but not as central drivers of AI policy nor as key sources of evidence or opinion. Since our research ended, we have noted (and participated in) much more work from civil society around new AI legislation.

Artists and arts communities, now some of the leading critics of generative AI, did not have a prominent place in our consultation documents. Such an omission is surprising as Canada has seen numerous debates about the future of creators in a digital economy as a result of the government's concurrent Broadcasting and Telecommunications Legislative Review. Perhaps we are still dealing with the split of Canada's former Department of Communication into two departments—ISED and Heritage Canada—as AI's economic impact seems to have overshadowed its cultural impact (Century, 2022). These concerns continue as Innovation, Science and Economic Development leads the Pan-Canadian AI Strategy with little connection, at least publicly, to Heritage or even those creators affected by the increased use of AI to manage and produce content.

Finally, a short note of self-reflection. Our list of silences did not include the ecological consequences of AI. As the environmental footprint of data centers and AI usage become contentious elsewhere, Canada has largely overlooked this issue. So profound is this silence that it was hard to place on our map as the environmental costs of AI are discussed only as a challenge for green computing projects (Hao, 2019).

Actors

Situational analysis helped us map all the actors and their discourses in Canadian AI governance. What we did not do, however, is identify the nonhuman actors involved in Canadian AI governance. Nonhuman actors could have included the land on which and towards which we should be having good relations. Instead, we identified more common actors.

Actor	Mentions
Government of Canada	8
Civil Society	7
Academia	6
Public	4
Vector Institute	3
Ontario Digital Service	3
Ministry of Finance, Ontario	3
MILA	3
Law Commission of Ontario	3
Gouvernement du Québec	3
Treasury Board of Canada Secretariat (TBS)	2
Statistics Canada	2
CIFAR	2
AMII	2
AI industry	2
Yoshua Bengio	1
Canadian Security Establishment (CSE)	1
Canadian provincial governments	1
Canada School of Public Service	1
Amnesty International	1
AI Impact Alliance (AiiA)	1
Access Now	1

Table 2. Key actors in policy documents

Classifying actors proved difficult, but we see two major patterns. Established players in the Pan-Canadian AI Strategy appear the most, as would be expected. The Vector Institute and Mila lead mentions along with government entities and all-encompassing categories, such as civil society, academia, and the public. There is a second gap marked by a clear underrepresentation of actors outside this small group; key firms and civil society do not appear among these major players.

The tight-knit ecosystem is particularly evident in federal efforts to consult the public on AI governance. Figure 3 maps the actors associated with the Government of Canada's *Open Dialogue: Artificial Intelligence in Canada in 2021*. Launched by the Public Awareness Working Group of the Advisory Council on AI, the dialogue involved only a small group and consisted of nineteen virtual workshops, mostly attended by university students, over the spring of 2021. The dialogue was a collaboration between the Canadian Institute for Advanced Research and the Université de Montréal, whose Algora Lab ran the dialogue. Algora is CIFAR AI Chair Marc-Antoine Dilhac's lab; he is also a member of the AI Advisory Council and participated in the Montréal Declaration. The consultations further marked an intersection between federal and provincial funding with the Quebec AI Observatory and IVADO both supporting the Algora Lab. The training guide for the dialogue does not mention human rights directly, but it does describe the Montréal Declaration as "ethical principles, not rules of governance, let alone legal norms. However, they can be translated into political language and interpreted in a legal manner, notably in the language of fundamental human rights".² The guide then introduces participants to a particular vision of AI governance.

² Available online at:

https://na.eventscld.com/file_uploads/3791cf7b32eb26534bc48f2724d017ee_GuideDeliberation_ENGpdf.pdf

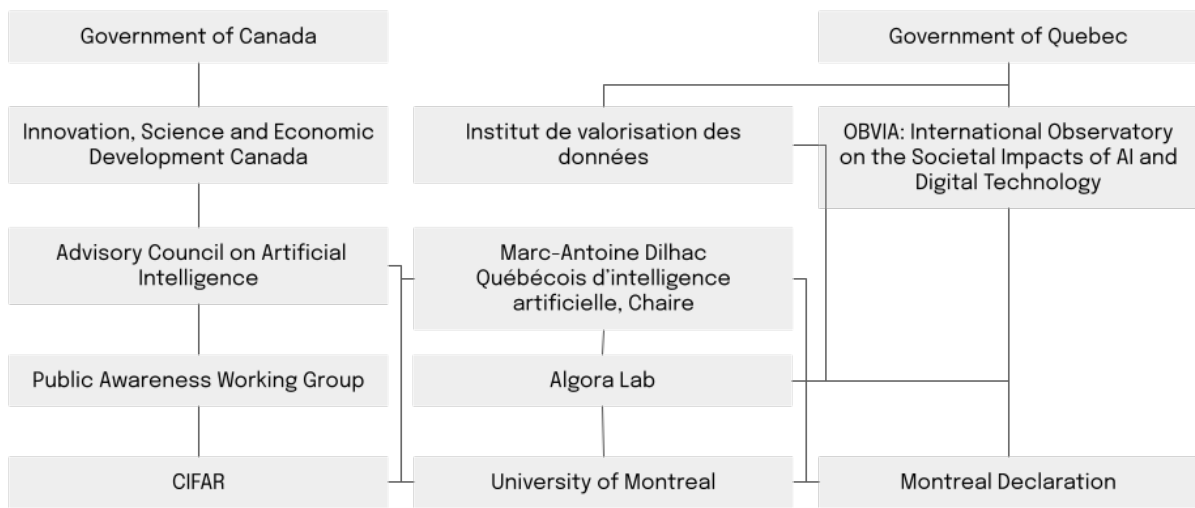


Figure 3. Insularity in public consultation

Such a closed group affects participation. The Public Awareness Working Group itself acknowledged as much:

future consultations must include more focused and deliberate outreach to under-consulted populations, including those most likely to be affected by AI technologies and systems. The lack of representation by these populations is a significant limitation of the Working Group's findings. The Working Group requested and received Council approval to continue their AI awareness work using a long-term and inclusive approach, that would specifically target participation from marginalized, under-represented, under-consulted, and Indigenous communities.(Government of Canada, 2022, p. 11)

The question remains, why were the consultations designed with such clear limitations in place from the start, drawing on the same pool of actors who were already direct beneficiaries of the AI funding system? What was the purpose of an “open dialogue” that enlisted largely the same actors? These concerns point to a deeper issue around consultation.

Conclusion

Our research concluded as the Trudeau government added a surprise addendum to its omnibus privacy bill in 2022: the Artificial Intelligence and Data Act (AIDA). AIDA sets out penalties and obligations for firms to ensure their applications do not cause individual harm. Novel in its shifting technology assessment from a matter of privacy to one of harms, AIDA ignored recommendations from the Office of the Privacy Commissioner, among others, to adopt a human-rights framework for artificial intelligence. Instead, as critic Bianca Wylie (2023) has pointed out, the influence of the economic imperative of AI is nearly total with discussions being led by Industry Canada. AIDA itself seems largely written for and by a nascent corporate AI consulting industry, looking for government regulation to legitimate its commercial approach to responsibility. Further evidence of this industry-friendly approach is provided by the ongoing efforts of the Standards Council of Canada to participate in drafting international AI standards in collaboration with established business standards organizations.

Our overall findings demonstrate that AI has arrived as a charismatic technology able to animate hopes and anxieties to renew Canada's fraught national project. From railways to pipelines, Canada's governments and elites have long relied on technologies and infrastructures to conjure the prospect of national unity. In many ways, Canada's national investments in AI continue longstanding efforts to justify confederation. This compelling narrative of Canada as a hub for "AI for good" is not a unique idea; however, it is easily replicated, allowing for a fragmented strategy all functionally meant to advance AI. We explore these questions further in our discussion of the structure of Canada's AI research system.

Canada's investment in AI has become a Faustian bargain. What we find is that the excitement around AI has also meant a growing reconsideration of governance through AI. Our introduction ends by calling into question the process of AI governance itself. A growing body of work asks whether AI is actually changing the nature of governance itself (Amoore, 2022; Fourcade & Gordon, 2020a). Joining a range of international scholars, we find that AI governance constitutes a shift in governmentality or, at least, reflects a push to reconsider governance through AI. We continue these themes in subsequent chapters on AI procurement, Canada's Algorithmic Impact Assessment, and attempts to integrate AI into Canadian immigration.

Beta governance was an enduring concept that developed through our team’s discussions. By “beta governance,” we mean that AI has often been framed as urgent and experimental. Here we draw on Gina Neff and David Stark (2004), who, over two decades ago, identified a trend toward “permanently beta” governance, where governance transpires always in a moment of testing and adjustment. We see this state of permanently beta as the norm in AI governance, sometimes rebranded under the name of “agile” governance. The opportunity here is for a more dynamic, participatory approach to Canada’s AI governance, but we find there has been insufficient innovation, as evidenced in our research on facial-recognition technologies consultations. Instead, beta governance is unstable, illegible, and inaccessible, allowing AI to be steered by players invested in and capable of keeping up with the rapid pace of change. Perhaps that might be a strategy for winning a global race, but the risk is that many will be left behind.

Chapter 2: Canada’s Algorithmic Impact Assessment by Nick Gertler

Author’s note: this chapter draws on my master’s thesis, see (Gertler, 2023).

How do you *do* “responsible AI?” Or “ethical AI?” Or “trustworthy AI?” Or whatever it is that we’re calling it these days. Consider the Montréal Declaration for a Responsible Development of Artificial Intelligence. Announced in 2017, the Montréal Declaration outlines ten principles for responsible AI: well-being, privacy and intimacy, solidarity, democratic participation, equity, diversity inclusion, respect for autonomy, prudence, responsibility, and sustainable development (Université de Montréal, 2018). To my ear, those principles seem sound; I too would like the development of advanced artificial intelligence systems to uphold the principles of equity, democratic participation, and so on. But how should those principles be upheld? The full text of the declaration offers some further guidance on what is meant by each principle, but it does not indicate how those principles should be put into practice.

In July 2017—four months before the Montréal Declaration was announced—Michael Karlin published a blog post titled “Responsible AI in the Government of Canada: A Sneak Peek.” Karlin worked for the Treasury Board of Canada Secretariat (TBS), the branch of the Government of Canada primarily responsible for overseeing the federal public service itself. The post brought together three related concepts: algorithms, AI, and automated decision making.³ Karlin outlined what he saw as both the potential and the risks for government use of artificial intelligence: government use of AI could mean more timely, efficient, and informed delivery of services, but poorly designed systems could ultimately undermine trust in government (Karlin, 2017).

³ For this project I use artificial intelligence and automated decision-making interchangeably. The TBS defines automated decision-making as, “Any technology that either assists or replaces the judgment of human decision-makers. These systems draw from fields like statistics, linguistics and computer science, and use techniques such as rules-based systems, regression, predictive analytics, machine learning, deep learning, and neural nets.” (Treasury Board of Canada Secretariat, 2023). For reasons I discuss, the kinds of technology assessed by an algorithmic impact assessment mattered less to this project as it developed. For more on the concept of algorithms, see: (Beer, 2017; Goffey, 2008; Zarsky, 2016).

In that same blog post, Karlin announced the development of a “Digital Disruption White Paper” on the use of AI in government. The white paper would inform the Government of Canada’s position on its own use of AI, and, exceptionally, it would be drafted in the open; anybody (but mostly experts in relevant fields) would be able to see and comment on the actual draft, of the actual white paper, as it was being written. Thus began the multi-year, quasi-open policymaking process that led to the development of the Government of Canada’s algorithmic impact assessment— the AI governance framework and software tool that is the focus of this chapter. This particular algorithmic impact assessment is significant because it is one of the first fully realized attempts at operationalizing the concept of “responsible AI.” Indeed, to the best of my knowledge, it is *the* first such attempt in a public-sector context.

Coming at a time when high-level AI ethics guidelines and statements of principle were being published left, right, and centre (Jobin et al., 2019b), Canada’s Treasury Board released its AIA as part of a concrete, mandatory policy for regulating government use of automated decision-making systems. In practical terms, the TBS AIA is a glorified online questionnaire that government departments should complete when developing an automated decision-making system (ADM). Respondents answer questions about the proposed system, such as, “Are clients in this line of business particularly vulnerable?” or “Will the system be making decisions or assessments that require judgement or discretion?” The TBS AIA is designed to be used in conjunction with the Treasury Board’s Directive on Automated Decision-Making. The directive mandates when departments must complete an algorithmic impact assessment and what impact-mitigating steps need to be taken depending on the AIA’s outcome. Most of the AIA’s questions contribute to a score corresponding to a series of risk levels in the directive. Each risk level comes with a specific set of requirements, such as whether the ADM system can render decisions without a human in the loop and whether (or to what extent) users of an ADM system must be informed of that fact.

My chapter acts as an explanation of an online resource, <https://aia.guide>, that serves as a guide to using and understanding the TBS AIA. The guide—which I will discuss in greater detail—is a collation of my background research for this project, presenting it in accessible language with an explicit focus on empowering users to think about how they could make the AIA their own.

The guide is also an attempt to deal with the encoded nature of the tool itself. By both describing and disrupting the logic of the AIA tool, the guide acts as a means of considering the ways in which the TBS AIA mediates AI governance.

My chapter explores the federal government’s algorithmic impact assessment (AIA) as what I call a policy media. The concept evokes how media arranges people and things in time and space, acting logistically (Peters, 2013). Rather than acting through norms or discourse, the AIA functions as code by requiring users to formulate policy, at least in its ideal use case, through interacting with the tool. Whereas policy documents or other policy instruments typically are encoded as text or rules, Canada’s AIA is something like an algorithm, technically an online questionnaire, that processes inputs into an output: a risk score. An update on the old Lawrence Lessig adage “code is law,” policy media instead eschews easy interpretation. Trying to study a GitHub repository seems straightforward but compiling all its facets can be overwhelming, so I instead tried a new method, building a website, to introduce and explain Canada’s AIA.

Situating Canada’s Algorithmic Impact Assessment

The Government of Canada may have been the first to implement an algorithmic impact assessment, but they were not the first to propose one. In the early stages of the tool’s development, the TBS AIA was never referred to as an AIA; that language came later in the development process, putting the TBS’s work more directly in conversation with proposals from academia and civil society. The two most notable proposals to this effect were from the AI Now Institute in the US and Nesta in the UK, both of whom released their proposals in February 2018—partway through the TBS AIA’s development (Karlin, 2018).

The Treasury Board’s questionnaire-based system for algorithmic impact assessment is quite different from other approaches. The model proposed by the AI Now Institute (AI Now Institute, 2018; Reisman et al., 2018) places a much greater emphasis on informing the public and communities about ADM systems that may impact them. The proposal’s authors underscore that algorithmic impact assessment’s core purpose is to engage those affected by ADMs. To that end, they do not set out specific assessment questions. Instead, they argue that

questions should be domain-specific and developed by particular departments. In the AIA developed by the TBS, every department and agency responds to a standard set of questions. Nesta's proposal—outlined in a blog post by Eddie Copeland (2018) titled “10 Principles for Public Sector Use of Algorithmic Decision Making”—more closely resembles the model that the Treasury Board eventually developed. The Nesta proposal never uses the term “algorithmic impact assessment.” However, principles like using a risk score to classify ADM systems and mandatory disclosure statements when users encounter ADM systems were directly taken up by the Treasury Board in the Directive on Automated Decision-Making.

The parts of Nesta's proposal that would require more government resources—either by way of additional development time or by requiring more funding—were not taken up by the TBS. Some aspects of what Nesta proposed, like the suggestion that governments develop an insurance scheme to compensate those unjustly affected by the ADM system, were simply outside the scope of the TBS's policy mandate; other parts of their proposal were adopted by the Treasury Board but in extremely limited ways. For example, Nesta proposed that a sandboxed version of the algorithm should be released for testing by external auditors. The TBS included a peer-review requirement in the directive, which is arguably a partial implementation of this principle. But it is implemented in such a way as to remove the open and public aspects, which are core to the principle in Copeland's proposal.

If we return to how responsible AI principles should be implemented, neither Nesta's nor AI Now's proposals directly answer that question. They do begin approaching an answer; unlike the Montréal Declaration, Nesta and AI Now considered their proposals' real-world implications and political feasibility. But, as I have noted, what makes the TBS AIA so interesting is that it is one of the first fully realized attempts at answering how responsible AI could be operationalized in a public-sector context. This is an impressive feat, but it came at a substantial cost.

Developing an AIA that (1) fit within the Treasury Board's preexisting mandate and (2) the public service could easily adopt became the central focus of the development process. As Michael Karlin put it when we spoke in 2022, “We wanted a lightweight algorithmic impact assessment that would take people no more than half an hour to do. So we wouldn't interrupt

their development flow too much, but still provide transparency” (M. Karlin, personal communication, March 18, 2022). The prioritization of creating something feasible and minimally disruptive, however, is also what led to many of the tool’s greatest weaknesses.

When comparing the Treasury Board’s AI governance framework with the AI Now Institute and Nesta proposals, the major omissions boil down to consultation and substantive public input. The TBS did not adopt the core elements of these proposals—direct engagement with affected people and communities—which would have introduced substantially more friction into the impact assessment process. These elements are undoubtedly messy, time-consuming, and more challenging to implement. Indeed, I would argue that many of the TBS AIA’s problems stem from this deep-rooted bias to action. The compromises that make the AIA easily digestible by its users within the public service have also been the primary sites of critique. As Moss et al. (Moss et al., 2021, p. 32) argue in their report on algorithmic impact assessment for the public interest, one of the core issues with the TBS AIA is that it fails to establish a meaningful “forum for accountability”. In the context of impact assessments more generally, a forum is a site that can “allocate responsibility for potential consequences of such systems and demand changes in their design, deployment, and operation” (Moss et al., 2021, p. 15). For a forum to provide a meaningful accountability mechanism, it needs to be able to require substantive changes to the system being assessed. But there is an inherent tension between creating a meaningful forum and the desire to create a lightweight tool that can be completed in half an hour.

As a result, the tool has not proven particularly effective at regulating government use of AI. While the directive is a mandatory policy instrument, it also features no meaningful enforcement mechanism. As a result, there have already been several high-profile instances where systems that should have required AIAs were developed and released without one (Cardoso & Curry, 2021; Keung, 2023). Moreover, even when an AIA is completed, there is ample reason to doubt whether a self-administered questionnaire can be a meaningful check on the development of complex algorithmic systems. Tellingly, the TBS AIA only faintly resembles other proposals for how to conduct algorithmic impact assessments.

One of the Treasury Board AIA's more distinctive features, however, is that it is open source. It was released by the Government of Canada under the highly permissive MIT License, and its source code is hosted in a public GitHub repository. As such, the TBS AIA is not a singular "thing." The fact that the Treasury Board made both the AIA tool and its questionnaire open source not only allows for its modification and repurposing in different contexts, but also encourages it. This feature of the AIA creates an opening, and my work on this project has been guided by an interest in exploring the potential created by the tool being open source.

Ultimately, this project is engaged with ongoing debates over the governance and regulation of AI. By deconstructing and disassembling TBS AIA, both outputs serve as methods for understanding and critically engaging a tool that lies at the heart of Canada's approach to regulating algorithmic systems. But *aia.guide* also offers a means of reassembling those pieces. In doing so, this project ultimately offers an approach to developing new participatory methods for AI governance and regulation.

(Re)Approaching the TBS AIA: An Overview of *aia.guide*

As I noted in the introduction, the guide is a website that acts as a resource for using and understanding the TBS AIA. The guide is divided into three modules: Platform, Policy, and Functions. The name of each module is intentionally ambiguous; they are meant to invite questions and encourage exploration rather than provide clear demarcations. Instead of reflecting existing divisions in the policy itself, each module offers a different approach to thinking about the TBS AIA system as a whole. In doing so, *aia.guide* seeks to both describe and disrupt the TBS AIA's internal logic. It requires users of the guide to approach the AIA as a (policy) media object rather than as a policy document, and it attempts to draw attention to the ways in which the TBS AIA system mediates AI governance.

The site has a table of contents in its menu for accessing specific topics, but the primary interface for navigating the site is a set of three large buttons—one for each module—placed after a brief introduction. To give users a jumping-off point, each module's navigation button is captioned with a guiding question that provides insight into the questions the module aims to answer.

Platform Module: What are the AIA's features as a policy platform?

The Platform module of aia.guide is primarily an attempt to defamiliarize the TBS AIA. Drawing on the formalist literary theory of Viktor Shklovsky, the concept of defamiliarization originally referred to the ways in which artistic and literary works can make everyday objects or experiences strange. Shklovsky argued that this process of estrangement could provoke a renewed appreciation of the everyday within the audience (Pangborn, 2010). Since Shklovsky initially identified defamiliarization as a literary technique, the concept has been rearticulated as a critical methodology. As Genevieve Bell, Mark Blythe, and Phoebe Sengers argue in their work on defamiliarizing the design of domestic technologies, defamiliarization as a method “calls into question our usual interpretations of everyday objects” (Bell et al., 2005, p. 154).

In this context, using defamiliarization as a method serves two purposes. On the one hand, it almost inherently makes understanding the tool more accessible to nontechnical users of the guide. The process of defamiliarizing the TBS AIA involves explicating each minor piece which comes together to form the AIA tool as a whole. As such, the level of prerequisite knowledge needed to understand the tool and its effects is substantially lower than when the tool is presented in its entirety. From a critical and analytical perspective, this same process facilitates (and arguably requires) a reevaluation of the AIA tool. In breaking the tool down this way, the guide requires users to consider what each piece actually does rather than defaulting to what we assume the system as a whole is doing.

In that spirit, the Platform module lays out the pieces underlying the software that the Treasury Board put together to host the AIA. It includes descriptions of what languages and frameworks the AIA tool is built on, a guide to the aia-eia-js GitHub repository and how to (attempt to) load a local copy of the AIA, and a walkthrough (Light et al., 2018) of the features and affordances of the questionnaire system. Many of these “features” would be considered mundane to anybody who has spent a reasonable amount of time on the internet. (Or anybody who has spent a reasonable amount of time filling out forms, for that matter.) I think it is fair to presume that most users of aia.guide are familiar with the difference between a checkbox and a radio button, for example. Nevertheless, the guide includes detailed descriptions of these elements and their contribution to the Treasury Board's AIA questionnaire system.

The ubiquity of these elements as standard pieces of the web is precisely the reason for this module. In the TBS AIA, each input type represents a design choice (and a policy choice) about what aspects of an algorithmic system can be evaluated. The most notable distinction to this effect is between scorable and non-scoreable inputs. Checkboxes and radio buttons might have different affordances, but ultimately they both produce quantifiable answers. Responses in free-text fields, on the other hand, cannot be straightforwardly quantified. As such, they *cannot* contribute to the score produced by the AIA. This score determines the system’s risk level and corresponding requirements under the directive; it is fundamental to how the AIA is meant to enforce certain practices. Free-text fields—which have proliferated in recent updates to the Treasury Board’s AIA questionnaire—do not contribute to this system.

Although they represent significant design choices in the AIA questionnaire’s development, the form elements’ mundanity makes them fade into the background. One’s focus tends to be on the questions being asked, not the inputs under them. Removing the question and presenting the inputs on their merits breaks a habitual and familiar usage pattern. In doing so, it draws attention to the inputs’ underlying significance and the general significance of software design to the AIA’s role in a policy assemblage.

Policy Module: How is the AIA used by the Government of Canada?

The guide’s Platform module sought to distinguish the TBS AIA’s approach to algorithmic impact assessment—the questionnaire and scoring model—from the policy context the Treasury Board deployed it within. That module focuses on the software object and its effects as a mediator. The specifics of *how* the Treasury Board and the Government of Canada use that platform are left to the Policy module.

Of *aia.guide*’s three modules, the Policy module is the most straightforwardly factual. It would be challenging to call this a “guide” to the TBS AIA without providing the context it was developed within, or how it is officially used. In this module, I largely attend to the role of the Treasury Board, the content of the TBS AIA questionnaire, and the Directive on Automated Decision-Making.

The challenge that the Policy module posed was not what content to include but rather how to present it. Certain subsections—particularly those on the AIA questionnaire and the AIA’s scoring system—follow the structure used by the Treasury Board in the directive and the descriptive landing page for the AIA tool. The main purpose of their inclusion in aia.guide was to bring all this information into one place. But describing the *processes* established by the directive was less straightforward. All the necessary information is in the directive, but it is both verbose and in extremely formal language. This may make sense in the context of an official policy document—especially when the expected audience is largely made up of public servants—but it does not make sense for a guide of this sort. I needed something shorter and more accessible.

Function Module: How does the AIA effect change?

The Function module of aia.guide centers the TBS AIA’s open-source nature. When I say “open-source nature,” I mean that quite intentionally. The fact that the tool is open source—the public availability of its source code—is not the focus of the Function module. Instead, the Function module attends to the *conceptual* implications of a policy being open source.

In this context, “function” refers to the *ways* that the TBS AIA can be used to effect change. In the guide, I liken it to a hypothesis—a theory of cause and effect. “Function” allows for an articulation of how we expect the policy to work, and it provides a falsifiable measure of whether it is living up to its intended purpose.

The utility of this approach to the TBS AIA only became clear to me after spending considerable time on this project trying to understand *how* this policy worked. This was hard precisely because the TBS AIA is riddled with contradictions: it is an enforcement tool without enforcement power; it purports to offer transparency through binary answers to vague questions; and it tries to provide public accountability, but only after the fact. These contradictions would not resolve themselves, and I realized that I needed a framework for understanding how they could all coexist.

Conclusion

The TBS AIA is a slippery object to study. Every time I felt I had a clear understanding of this strange policy object, a new fact or facet would emerge and upend everything. As I came to realize, developing a complete and accurate history of this tool is a near-impossible task. The quasi-open development policy process that produced the TBS AIA suggests that a tempting archive of the policy's development might exist. But in practice, the archive is scattered and incomplete; once-public Google Docs have been made private, other files have been taken offline, and the Government of Canada's secretive internal culture puts key non-public documents out of reach.

As Tess Lea reminds us in *Wild Policy*, though, this teleological approach to policy misses the messy bigger picture (Lea, 2020b). Not only is writing a biography of this particular policy and tool extremely challenging, but it may not be desirable in the first place. Through interviews, access to information requests, and internet sleuthing, I have reconstructed a “good enough” sense of the TBS AIA's history. But these are pieces of a fairly incoherent policy assemblage. It would be dishonest to present them as a complete picture.

In the end, I am informed by a deep ambivalence over the technology (Bucher, 2019). This AIA tool is largely performative, yet its significance marks an important tactical intervention. My project has tried to study the AIA through a critical guide designed to use and reflect on the tool's potential, as well as new possibilities for AI governance. There are undoubtedly critiques to be made of the TBS AIA. Indeed, this project has further developed several of them. Attempting to assess algorithmic impact using a one-size-fits-all, self-administered questionnaire is a flawed approach. And, as I have demonstrated, certain decisions like including additional unscored questions in more recent versions of the questionnaire have made it all the less effective. The tool's use by the Government of Canada remains deeply questionable and largely performative.

Despite these problems, however, I have still chosen not to write a polemic against this policy. Perhaps the key contribution of the TBS AIA—and indeed the thing that drew me to it in the first place—is that it was an early attempt at putting principles into practice. It certainly would have been easy enough to write a purely critical evaluation of this tool, but doing so would not have produced many answers about where to go next. By exploring the TBS AIA through

research-creation, it has been possible to develop a deeper critical understanding of the tool while also trying to find a path forward.

As for the AIA itself, the place where I see it as having the most potential is as a platform *for* algorithmic impact assessments. The questionnaire approach can be stifling when it is top-down, but it shows tremendous promise as a model for bottom-up consultation processes. In this way, reorienting the TBS AIA shows how we might—in practice—begin to develop inclusive models of AI governance that are not de facto driven by experts. Whether or not those models ultimately make use of this particular tool, hacking the TBS AIA gives us an avenue toward participatory methods for governing AI systems.

Chapter 3: AI Source List: The Ethics of Procurement, the Display of Responsibility, and the Production of Nonknowledge by Meaghan Wester

This chapter draws on Meaghan Wester’s master thesis, see (Wester, 2023).

Introduction

Palantir is a predictive analytics services company. Their clients include governments across the globe as well as law enforcement and Wall Street firms. Their name—a reference to a fantasy crystal ball—and contracts suggest the company can be eerily omnipresent. Peretti writes, “Palantir does not just provide the Pentagon with a machine for global surveillance and the data-efficient fighting of war, it runs Wall Street” (Peretti, 2017). Beyond the company’s presence in influential sites of power in the public and private sectors, their services have played a role in human rights–threatening and infringement cases in Iraq, have been used in predictive policing in the US, and have aided US Immigration and Customs Enforcement (ICE) during migrant raids (The Guardian, 2017; Amnesty Report 2020). Despite these controversial uses, Palantir remains qualified to supply AI to the government of Canada. Murad Hemmadi (2019) at the *Logic* points to a paradox in the regulation of AI suppliers such as Palantir. As he highlights, Palantir is on Canada’s responsible AI source list of preapproved suppliers: “But officials in Ottawa weren’t able to consider the company’s most controversial work when evaluating whether it should be allowed to bid on some Canadian government artificial intelligence projects”. If Palantir has an international reputation as a contentious technology firm, how did it evade controversy while becoming approved as an ethical AI vendor to the government of Canada?

This chapter uses the case of Palantir to understand the intersection between ethics, procurement, and Canada’s AI source list of preapproved suppliers. This list provides a prime site to study AI governance as it is both (1) the mechanism by which government organizations procure AI and (2) how Canada claims to enforce its normative framework for AI (i.e., responsible AI). The AI source list’s requirements, M1 and M2, ask interested

suppliers to demonstrate their expertise and capacity, while M3 explicitly requires suppliers to outline how they will behave ethically and responsibly in delivering AI tools and services. As such, M3 explicitly situates private actors' ethical stances as part of the procurement process. Procurement surfaces the administrative processes by which AI embeds in and modifies public service. But analyzing procurement also allows us to study the limits of the power the government claims over AI providers, procurement purchases, and regulating AI more broadly.

What this chapter finds is that procurement is not merely a narrow policy studies topic but a locus where the political theory of AI governance can be studied—indeed, procurement should be understood as central to political theories of AI governance. Traditionally, procurement is a matter of good governance, fitting into another routine application of the policy sciences. But procurement is better understood as a critical nexus of debate about the future of governmentality. In fact, this logic is the result of the New Public Management (NMP) paradigm that is integral to the neoliberal turn of the 1980s and emphasizes regulatory liberalism, privatization, and market-based solutions (A. Clarke, 2019; Glor, 2001). The NMP was directly imported from the business world to the public sector and is a way of doing politics otherwise without appearing to do so. Reframing citizens as consumers, the imperative of efficiency, and restructuring management to mirror corporate structures all helped create a brittle public administration that strategically must rely on the private sector to deliver services to citizen-consumers. Procurement, under this logic, is a critical subject at the intersection of state-corporate relations. Changes in procurement, thereby, offers an important litmus test to gauge contemporary Canadian and international trends in governmentality.

Given the significance of procurement in thinking through shifts in governmentality, I find that in the case of the AI source list, it functions as a site of nonknowledge. The loosely related fields of the sociology of ignorance, ignorance studies, and agnotology are concerned with varying types of absences of knowledge and how they relate to epistemology, power, and knowledge. Intentional ignorance, amnesia, unknown unknowns—these are some of the many forms of nonknowledge of concern for these interdisciplinary fields. When applied to bureaucratic processes and policy, the study of nonknowledge is often mobilized to suggest

that *more* knowledge must be uncovered or produced to palliate emergent policy needs. But applying nonknowledge to policy studies has outlined how routine forms of institutional nonknowledge, ignorance, and amnesia coexist with the emergent complexity of governing through unknown unknowns, such as during the COVID pandemic

(Hannah et al., 2023; Parviainen, 2020; K. T. Paul et al., 2022). These applications of the concept paint a more nuanced picture and create space for studying nonknowledge as productive and, in some cases, strategic, as in the control of forest fires (Kuchinskaya, 2014). Beyond a conceptual tool to think through productive sites of absence of knowledge, nonknowledge's relationship with incalculable risks makes it a prescient concept to unpack the nuances of AI governance, which is ubiquitously framed as a risk management issue (Beck, 2008; Gross, 2016; Wehling, 2011).

Drawing on the fields of ignorance studies and agnotology in STS, this chapter considers how nonknowledge is itself a governance mechanism and how it works to foreclose how AI ought or could be governed. I read nonknowledge applied to policy studies as closely related to the concepts of negative policy and silences in policies (Freedman, 2010). Through this lens, the AI source list is strategic in carving out areas of unknowability at the site of M3. Understanding the AI source list as a productive site of nonknowledge reconciles how the list can at once require suppliers to outline their ethical approaches while at the same time not making sense of them.

The case of Palantir shines a light on what the list fails to capture and what it achieves. Findings suggest that the list merely fast-tracks suppliers without robust assessment of responsible practice. In reviewing the suppliers' evaluation grids, I was surprised to see that at almost every quarterly update of the AI source list, some suppliers failed this M3 criterion. In most cases, suppliers who failed to meet this requirement wrote that "AI ethics did not apply to them because of the type of service they delivered." The reviewers' attention to detail and expertise, more than the M3 requirement itself, form the grounds for approval or rejection. The main reason why the list fails to do so is that rather than serve as a site where governments impose explicit requirements on prospective suppliers, it is a space where nonknowledge is produced. As a result, multiple suppliers qualified for the list without robustly demonstrating a clear definition of what it means to be responsible with AI. The preapproved AI suppliers

source list is thus at once a site where responsibility is displayed and nonknowledge is produced.

What Is Procurement? What Is Prequalification?

In Canada, public procurement does not fall under a centralized federal jurisdiction, which means that the different procurement processes are dealt with internally by departments and agencies (Fasken, 2020). As a result, policies are scattered and localized within different government silos. Each department oversees the internal approval process through directives and policies. These directives and policies can change during a single tendering period, which makes the process highly flexible and uncertain for suppliers. Some departments are, however, constrained by trade agreement requirements that can provide longer-lasting indications of requirements. The AI source list is a product of the modernization of Public Service and Procurement Canada (PSPC) and represents an attempt to deal with the siloed tendering of AI.

Public procurement started modernizing its process in 2018 through the Electronic Procurement Solution (EPS) and the Buy and Sell platform, transitioning gradually to the CanadaBuys platform. The 2018 Budget allocated “\$196.8 million over 5 years to establish an e-platform for simpler, better procurement” (PSPC, 2021). PSPC (2021) explains, “Canadian companies have long asked the federal government to improve its relationship with suppliers by making opportunities easier to find, simpler to navigate and faster to award, with less administrative burden. Government procurement to date has been heavily paper-based with limited self-serve options for suppliers”. The bidding process to provide PSPC with IT solutions to modernize itself is a different one than the AI source list. They both, however, follow an Invitation to Qualify (ITQ) model to prequalification and promote “agile” procurement, and many suppliers who qualify are on both lists.

The Government of Canada has been trying to modernize its procurement. In 2018, as PSPC and the Treasury Board of Canada Secretariat worked on this project, they could choose from myriad mechanisms to organize public procurement of AI. As outlined on the AI source list page, “Public Services and Procurement Canada (PSPC), together with the Treasury Board of Canada Secretariat (TBS), held a procurement process to establish a list of suppliers who can

provide the Government of Canada with responsible and effective AI services, solutions and products” (Treasury Board Secretariat, 2020). In this move towards AI governance, Canada has aligned themselves with the Canadian brand of “responsible” rather than “ethical” AI. In the absence of any formal or legal definition of “responsible AI,” the list, among other initiatives, becomes a key site to observe how Canada understands and operationalizes the idea. The mechanism of prequalification thus serves the double function of AI governance and public procurement of AI.

While tracking down how they arrived at this decision is beyond the scope of this report, it is possible to outline critical reasons why the prequalification mechanism was deemed adequate for the list. Berjis (2012) defines prequalification as a process “to assess the capability and competence of potential bidders through screening of contractors according to a given set of criteria” (Berjis, 2012, p. 1). In Canada, prequalification involves the following steps:

1. Vendors submit responses to specific requirements to the Invitation to Qualify (ITQ).
2. Their submissions are assessed by the appointed PSPC committee.
3. Vendors who qualify are subsequently invited to bid on specific tenders.

These steps contrast with the more traditional and single-layered tendering process where a Notice to Public Procurement (NPP)—a specific tendering bid containing the contract details and relevant departments’ requirements and demands—is published. In that case, interested bidders submit bids directly to the NPP, which are then evaluated by PSPC based on the client department’s stated requirements. Prequalification necessitates an additional step and takes place prior to a specific NPP being published. Under prequalification, interested suppliers submit responses to the Invitation to Qualify (ITQ), which consists of responding not to specific bids but to requirements; in the case of the AI source list these are M1, M2, and M3. This pre-tendering step presents distinct advantages and limitations for the government as well as for the sellers.

If suppliers choose the prequalification route, the steps from their perspective consist of the following. The ITQ can be found through the Buy and Sell platform and the GC platform;

multiple business blogs and trade organizations also circulate ITQs and NPPs. Companies submit their responses to M1, M2, and M3. A committee evaluates their submissions according to already established requirements. (ATIP A-2021-00322, p. 1). The vendors are then classified into bands according to the project budget. The departments choose three suppliers they would prefer to compete for their bid, and PSPC randomly selects seven others for a total of ten (Treasury Board of Canada, 2021). Lastly, the winning supplier is notified and receives the contract.

A vendor may choose not to go through prequalification based on how confident they are and how many resources they have to respond to a specific department tender notice. In addition, the industry within which the vendor competes may also impact this choice as tenders that fall under trade agreements tend to be more static, which makes it easier for the vendors to keep track of the otherwise often-changing requirements (Fasken, 2020). In short, the list is a mechanism to promote agile procurement—however, it performs much more than this work.

Displays of Responsibility

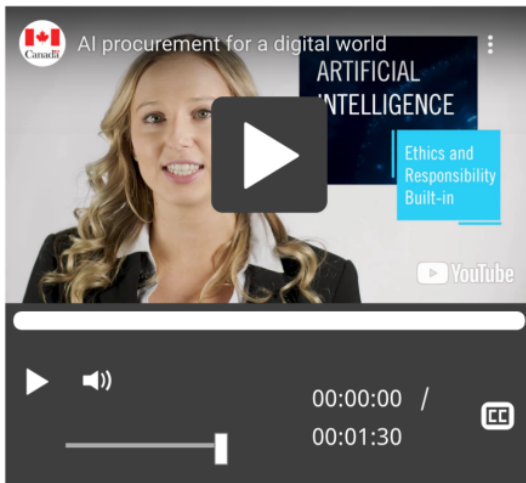
AI strategies characteristically exist in multiple formats, documents, and locations and are subject to change (Bareis & Katzenbach, 2021b). In centering “responsible AI,” the government of Canada also deploys the fluidity characteristic to AI strategies that makes the assessment of and public critical engagement with the strategy nearly impossible given its spread, iterative character, and necessary partiality. One site where elements of this strategy coalesce is the government of Canada’s webpage that outlines its twin approach to the responsible use of AI initiated in 2018. The page presents the Algorithmic Impact Assessment (AIA) directive and the AI source list as a twofold governance approach. Given the flexibility of AI strategies, this landing page affixes the centrality of procurement and AIAs to Canada’s responsible AI approach, and it, crucially, also produces a strategic and partial form of visibility for this approach.

Responsible use of artificial intelligence (AI)

Exploring the future of responsible AI in government

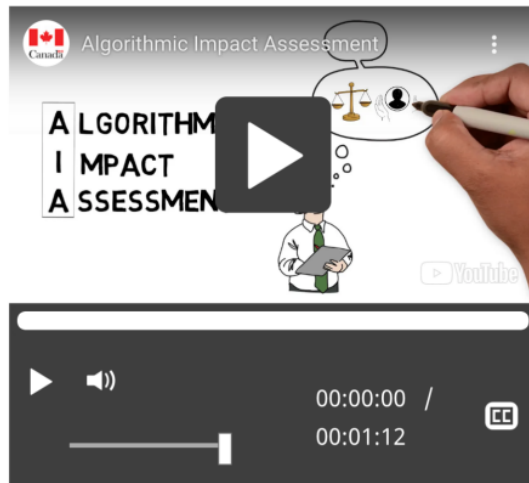
Artificial intelligence (AI) technologies offer promise for improving how the Government of Canada serves Canadians. As we explore the use of AI in government programs and services, we are ensuring it is governed by clear values, ethics, and laws.

AI procurement for a digital world



▶ AI procurement for a digital world - Transcript

Algorithmic Impact Assessment



▶ Algorithmic Impact Assessment - Transcript

Figure 1. (Government of Canada, 2023) The figure is a screen capture of the government of Canada’s website on the responsible use of AI.

Much like a display cabinet in a shop or a museum, the government’s responsible AI page is a strategic display of transparency. Importantly, as the map displays (see Wester, 2023), the AI source list is a joint project between TBS and PSPC. Two videos explain the federal responsible

AI strategy's two initiatives. The video titled "Artificial Intelligence: Ethics and Responsibility Built-in" describes the role of the AI source list; the narrator states:

While AI is a powerful tool, it must be used responsibly. We have to eliminate bias, be open about how AI is informing decisions, and ensure potential benefits are weighed against unintended results. That is why we build responsible use into everything we do, including our first AI procurement process.(Treasury Board of Canada, 2021)

This video, along with the AI source list landing page and other documentation of the strategy, acts as a vitrine for displaying responsibility. Crucial to this partial visibility is that responsibility is not defined, benchmarked, or enforced; instead, it is used in a display context. The list promises to simplify procurement for the departments by grouping suppliers in one location and accelerating the process. However, departments need not buy only from those on the list (Public Gathering #2, 2022). TBS highlights that "many [vendors] are proud to be on the list and use it in their publicity. Beyond this, it is a potential source of contracts for them" (Public Gathering #2, 2022). Accordingly, if the list is used, the requirements can enforce "responsibility"; if it is not used, the requirements merely display "responsibility."

As of April 2023, there are 117 companies registered as part of this list, ranging from consultancy firms (such as Accenture, CGI, Deloitte, KPMG), universities (such as Simon Fraser University), Big Tech (such as Amazon and Microsoft), and military and defense contractors like Palantir. Companies that do not submit themselves to the prequalification process can also do so to avoid being subjected to this (light) scrutiny. Consider how Clearview AI and Palantir initially offered their services pro-bono, effectively hacking the procurement process (ATIP A-2020-00060, n.d.). Both companies offered free trials to different government organizations for specific instances (respectively, various law enforcement agencies and the Public Health Agency of Canada), thereby evading or bypassing the prequalification requirements as well as the more general tendering steps outside of the list.

For a company such as Palantir that has featured in multiple media controversies but whose main clients are governments, the list is a prime place to show that they are willing to play by the rules. What is less obvious when being presented with the vitrine is how lax these rules are.

Site of Nonknowledge, M3, and the Ethics of Procurement

Every supplier needs to provide an Ethics Statement, or M3. M3s are one of the three requirements suppliers must submit to qualify for the AI source list. The M3 requirement states: “The supplier must provide examples of how it addresses ethical* practices when delivering AI. This includes demonstrating experience in applying frameworks, methods, guidelines or assessment tools to test datasets and outcomes” (ITQ EN578-180001/A, 2018).

While laudable in spirit, the M3 is in actuality a box-checking exercise. Through material acquired through an Access to Information and Privacy (ATIP) request, I was able to review the grid used by the evaluation committee to assess suppliers’ submission to the ITQ (see Wester, 2023). The evaluation is simply a binary choice between whether ethics have been met (1) or not met (0). There is no further consideration of the normative stance taken by the company—to demonstrate that the supplier is ethical is sufficient.

In reviewing the suppliers’ evaluation grids, I was surprised to see that at almost every quarterly update of the AI source list, some suppliers failed this M3 criterion. In most cases, suppliers who failed wrote that AI ethics did not apply to them due to the type of service they delivered. No suppliers were rejected on the grounds that their ethics did not align with the normative aims of TBS, only on the grounds of not or insufficiently considering “ethics” abstractly.

Examples of justifications for companies who met the M3 requirement highlight that the reviewers’ attention to detail and expertise, more than the M3 requirement itself, form the grounds for approval or rejection. One reviewer noted in the rubric assessment of one company that met the M3 requirement: “Company principles are high-level, but cover important components of quality, bias, choosing uses (benefits society, not weapons), focus on testing, privacy and human rights” (p.1499). Another reviewer wrote in another requirement-meeting company’s review: “They have clearly articulated ethical principles on page 12 that

align well with the ethical framework established by TBS. (i.e. human in the loop, open source, etc.)” (p.1505). Both instances justify that the requirement was met based on the mere presence of a loose definition of ethics.

The M3 mechanism fails to actively regulate as it does not consider what AI providers normatively propose in these submissions. In effect, M3s can only evaluate if there is an ethic considered. Matching TBS/PSPC best practices on data and ethics is binary: met or not met. The type of ethic or the reasoning why these practices are in place remains undifferentiated in the evaluation grids. As such, the M3 criteria strategically produces nonknowledge of what the suppliers stand for ethically and, in turn, lets the list off the hook of having to clearly state cumbersome ethical requirements.

This “met (1)/ not met (0)” framework does not capture the two tensions in governing AI surfaced by situational mapping (see chapter appendix): (1) where power is located and how it is distributed in governing technology (compliance versus self-rule) and (2) ethics of technology stances regarding whether technology is neutral or agential. The failure to assess suppliers’ positions on these issues makes it impossible to engage with their normative stances critically.

The vast variability between submissions is indicative of the evaluation protocol’s inability to assess AI providers’ stances on how the power to govern AI ought to be distributed. Even if all twelve suppliers in the sample met the M3 requirement, they all hold different views on how and where the authority and legitimacy to govern AI is located between the private and public sector. Some suppliers barely comply and meet the values listed under M3; others use their M3 submission to form recommendations.

Three suppliers in the sample—Ova Inc., Donna Cona, and Chillwall AI—take a stance close to compliance. All three submitted very brief, bare bones submission that strictly followed the guidelines’ keyword suggestions. Ova Inc. emphasized their human-in-the-loop or human-on-the-loop approach to achieve explainability, pointed to their use of alphanumeric codes replacing participants’ names in lieu of privacy guardrails, and highlighted their fulfillment of the GBA+ (Gender-based Analysis Plus) Canadian guidelines with their 25% of their workforce consisting of women. Needless to say: these measures are rudimentary. Donna

Cona, in a joint venture with MasterTech, and SFU both stayed close to the language in the guidelines and kept their M3 brief. Chillwall AI submitted less than a page for their M3, which contrasts starkly with Deloitte's 83 pages and Palantir's 52. Chillwall, who sell emotion-detection technology, attempt to dispel privacy concerns in their M3 by stating: "Events data that we use is non-personal and thus there are no questions or concerns that could be raised with regards to privacy" (Chillwall AI M3, p.6). These examples highlight how the line for "met" or "compliance" does not require suppliers to understand or demonstrate the breadth and depth of their ethical and political role as a supplier of AI to the public sector.

We can learn more about nonknowledge production in the M3 through Palantir's submission. Palantir is one of the suppliers who submitted an M3 with thick ethical recommendations, they write:

We believe that context matters and is essential to determining whether specific types of information should be used in AI. For example, it is one thing to seek to analyze whether sensitive categories such as race and gender linked to genetic attributes correlate with the prevalence of certain diseases. It is quite another thing to seek an association between the same immutable personal attributes and one's likelihood to default on a loan. (Palantir M3, 2018, p.18)

The dissonance between this passage and the way Palantir's technology is deployed in the world (e.g., its use by ICE at the US-Mexico border) saliently exemplifies what the current way to discuss and measure AI ethics fails to account for: what a supplier makes a commitment to and how they act can be vastly different. The promise to regulate AI through procurement and the prequalified list fails so long as suppliers retain the power to self-govern. Further, the list produces nonknowledge by hiding the differences between ethical stances and self-governance approaches from the Canadian public while facilitating the incorporation of these systems into public administration.

Conclusion

The AI source list fast-tracks public procurement of AI and displays responsibility; it does not enforce responsible practice and instead produces nonknowledge. This is how companies like Palantir qualify to be on the preapproved responsible AI suppliers list. The Pre-Approved Suppliers List is a site of production of nonknowledge that solicits ethical statements about AI but crucially does not entail knowing how to read or interpret these statements. Intentionally or not, the unknowability of what is an ethically sound ethics statement allows anything submitted to gain approval. In light of this research, I echo calls to complicate the relationship between knowledge and ignorance as it relates to governance and governmentality and consider what is purposefully unknowable.

What does this mean more broadly for the political theory of AI governance? I explore this shift in governmentality through what I call the Procurement Circuit. The relationship is circular as private companies' structure how the public services approach AI just as the Canadian public services try to define their approach to AI when procuring it. Conversely, the government attempts to structure how AI providers manufacture AI (i.e., responsibly) through procurement guidelines and other policies. As clients and regulators of AI providers, governments do not merely govern AI they also govern with and through it. The procurement circuit refers to the co-shaping of governance through procurement. The circuit produces two interlocking problems. First, it leaves AI providers' legitimacy and authority over normative questions on AI unquestioned. Second, it prevents Canadians from assessing, debating, and resisting the aims of policies regulating AI and through AI. As AI systems restructure internal processes and policy making, AI suppliers become entrenched in public administration. This chapter underscores the asymmetry of power between the government and AI suppliers within the procurement circuit.

Thus, governing AI at the site of procurement is crucial, but it currently acts more as a site of nonknowledge that forecloses how AI ought or could be governed. The display of "responsible" and production of nonknowledge lets the government off the hook for establishing solid standards for selling AI to the public sector. Instead, AI ethics being deemed uncontroversial contributes to the closure that takes place at the site of the AI source list. Indeed, the fact that Palantir has ethical clearance when its technologies have been used for

violent means raises the notion that perhaps AI ethics has become abstract and unconnected to the real world, which enables it to be so uncontroversial. More broadly, the use of responsible AI and the list deceptively resolves the question of what technology is good for society, thereby foreclosing democratic debate while legitimating private tech firms as ethical and political leaders in public life, services, and administration.

Appendix: Methods

Evidence for this chapter came from a series of Access to Information Requests (ATIPs). I relied on Access to Information Requests to access two types of documents: 12 suppliers' submissions to the M3 requirement as well as the evaluation rubric the approval committee files to evaluate them. To supplement these rubrics, I also requested the training material evaluators receive. Out of the 117 suppliers on the Pre-Approved Source List, I narrowed it down to 12 suppliers across the different service categories, size, and amount they had been preapproved for. The necessity of narrowing the list stemmed from interactions with our assigned ATIP officer, who recommended, given the time constraint of the study, to request a more concise list. The 12 were Deloitte, Palantir, IBM, Element AI, Simon Fraser University, Donna Cona in joint venture with Mastech, KPMG, Chillwall AI, Service Now, PricewaterhouseCoopers, Ova, and Thales.

The ATIP material was used to produce two conceptual maps drawing on situational analysis (Clarke et al., 2017). Situational analysis (SA) offers three main cartographic approaches building upon and extending Strauss's situation-centered "social worlds/arenas/negotiations" framework (Clarke, 2005, p.xxii). As part of my MA thesis, I produced two maps: the first outlined how suppliers and interested departments arrived at the AI source list through organizational pathways. The second is a map of the ethical stances taken by suppliers in their submissions. To create this map, which normatively situates the suppliers, I inductively coded a total of 410 pages (including redacted pages) of M3 submissions from the 12 suppliers and overlaid the coded documents over a literature review of four areas in the ethics of technology. The coding criteria include governance and AI narratives, Fairness-Accountability-Trustworthy-Explainable (FATE) keywords, definition, or explanations as well as hyperlinks

directing reviewers toward external material or achievements. The normative maps emphasize how suppliers' answers differ along two key axes: (1) where power is located and how it is distributed in governing technology (compliance versus self-ruling) and (2) ethics of technology stances on whether technology is neutral or agential. For instance, Palantir is a particularly interesting case because their submission to the M3 requirement is very detailed and demonstrates how suppliers understand themselves as providers of thick ethical standards. They therefore align more closely with other large players, such as IBM and KPMG, who understand their role to extend beyond compliance and into self-regulation and their submission as a space to assert themselves as leaders in governance. A passage in their M3 that illustrates where Palantir stands on the first axis states: "Beyond what is legally permissible, an ethical framework is critical to establish boundaries when navigating the complexities of novel technology applications, such as artificial intelligence" (Palantir M3, 2018, p.17). This chapter draws on findings that emerged from both maps.

Chapter 4: Consultations on Facial Recognition Technologies by Maurice Jones and Fenwick McKelvey

It's dangerous, racializing, and has few legitimate uses; facial recognition needs regulation and control on par with nuclear waste (Stark, 2019).

On February 13, 2020, the Toronto Police Services (TPS) admitted that some of its members had used facial recognition technology (FRT). Their confession came a month after denying allegations of using an application developed by the controversial USA-based company called Clearview AI. The company created and maintained a database of more than three billion images scraped from the internet, including Facebook, YouTube and Venmo, without people's consent. It then compiled the images into a proprietary image database called a "Face Library", which it sold to law enforcement, immigration services, universities, high schools and businesses. The database enabled it to match photographs of people against the images in the database using facial recognition technology. The pressure on the TPS and other law enforcement agencies came in the wake of a New York Times article that broke the news that more than 600 law enforcement agencies across the United States and Canada used FRTs (Hill, 2020). Canadian law enforcement agencies, including the TPS, were out of the frying pan, into the fire when on February 26, 2020, Clearview AI revealed it had suffered a data breach and its entire list of customers was made public.

The scandal sparked widespread outcry in the Canadian media since it revealed that police services across the country used facial recognition systems that were based on illegally and secretly collected data, in clear violation of federal and provincial privacy laws, by scraping images from the internet without permission. These developments put pressure on policy makers to address the issues raised by FRTs and similar AI-driven tools in policing. The TPS, the Office of the Privacy Commissioner of Canada (OPC) and its provincial counterparts, and the parliamentary Standing Committee on Access to Information, Privacy and Ethics (ETHI)

were at the forefront of this policy response, launching investigations and calling for policy consultations.

Clearview AI makes a fitting introduction to our report about FRTs and consultations as a key tool to address emerging controversies around AI. In this report we ask the following three questions: 1) Which actors are enrolled and which are silenced by policy consultations? 2) How do policy consultations enroll actors and construct issues surrounding FRTs? 3) How do consultations affect concrete policy outcomes?

This report focuses on four policy consultations that were held on FRTs in Canada following the Clearview AI scandal. These are the two consultations held by the Office of the Privacy Commissioner (OPC), the consultation by the Toronto Police Service Board (TPSB), and the hearings of the parliamentary ETHI committee. Highlighting these four consultations aids in understanding how different forms of consultations enroll different actors and construct issues differently. In addition, it tells if and how public opinion and outrage translate into policy recommendations. In addition, our team of researchers initiated an independent policy consultation organized through an online survey with civil society actors. This survey was focused on understanding whether civil society actors got enrolled, reached out to, or silenced by consulting parties, as well as these actors' understanding and opinions on FRTs.

Our analysis finds that the four instances of consultations enrolled a notably different set of actors depending on the consulting party, the specific framing of the topic, the forms of participation (such as written or in-person), the length of the consultation period, and their media publicization or lack thereof. Most notably, consultations became powerful closure mechanisms for controversies surrounding FRTs and AI; mechanisms that tend to sideline raised concerns and foreclose any sort of debate. The most powerful example is that a diverse range of participants to all four consultations raised the need for an outright ban and/or a moratorium on FRTs and similar AI systems until adequate governance measures are in place. All but one of the resulting policy documents, the report by the ETHI committee, took up this fundamental concern to recommend a temporary moratorium on FRTs. Ultimately, we conclude that consultations are merely one among many factors that play into the governance of FRTs and AI in Canada, with institutional legacies, agenda setting by powerful individuals,

partisan politics, and the simple need to “appear to be doing something” equally driving the governance process.

In the following paragraphs we set out with a brief overview of the underlying methodology for this report. Following this, we present a timeline of how FRTs in policing emerged as a controversy in Canadian AI governance. The next four sections present an overview of the four consultations leading up to our analysis and, finally, concluding remarks.

Methodology

The methodology for this section of the report centered around situational analysis as put forward by (A. E. Clarke et al., 2018). Situational analysis is specifically attuned to mapping actors, issues, and their relationships, which are rallied by policy consultations surrounding FRTs in Canada.

We set out with an initial survey of policy documents related to the four policy consultations on FRTs we address in this report. These included public calls for consultation, draft policies posted for consultation, policy documents resulting out of consultations, What We Heard reports, and publicly available submissions. Submissions to the two consultations run by the OPC were not published due to privacy and translation concerns. Access to information requests to gain insights into these submissions are currently underway.

In addition to policy documents, we surveyed media reports to establish a timeline of events with a specific focus on the emerging Clearview AI controversy and to further contextualize consultations within larger processes of AI governance in Canada. This survey of policy documents and media reports gave us an initial understanding of the timeline of events, on the consulting parties, the framing of issues posted for consultation, and the actors that contributed to consultations.

In order to gain deeper insights into how policy consultations are organized, how actors are enrolled, and how issues are constructed, we conducted a total of eight interviews with nine interview partners representing policy makers, politicians, law enforcement, civil society, and journalists. These actors were either directly involved in the consultations as organizing or

participating parties and/or had additional expertise on the topic through their professional investment. The interviews were conducted online and recorded via the Zoom platform for a duration of 60 minutes.

While the initial mapping showed us who was present in the consultations, it said little about which are the voices that might be absent or silenced by these processes, and more importantly why that might be the case. In order to understand these, we conducted our own online survey targeted at civil society actors, most notably those representing marginalized communities that are most affected by the deployment of FRTs in policing.

FRTs in policing: a controversy a long time in the making

FRTs arrived in Canada well before these technologies became a matter of public concern. Since the early 2010s, police forces across Canada have experimented with predictive policing, algorithmic risk assessment, and facial recognition (Robertson et al., 2020). In Toronto, the TPS started their own trials with AI and FRT-adjacent technologies in 2018 when the city purchased a mass surveillance technology, ShotSpotter (Ezeonu, 2009; Gillis, 2018). The same year, the TPS offered a \$450,000 contract to NEC's Intellibook for a "Facial Recognition System." The contract ended a year-long procurement. With little public disclosure, the TPS connected its mugshot system with a facial recognition tool. FRTs did become an issue in some Canadian municipalities following similar bans in the United States. The Montréal city council voiced concerns and initiated an inquiry into the use of FRTs by police forces as early as August 2019 (Harris, 2019; Serebrin, 2020). None of these developments prompted consultations.

Clearview AI started two of the consultations considered, though not at first. News of Clearview AI was not initially controversial in Canada. The story first broke on 18 January 2020 in the New York Times. Kashmir Hill's story warned that the American firm had trained its own facial recognition system using billions of images uploaded onto social media. Hill mentioned that Canadian law enforcement had used the tool as well, without specifying which branches. A few days later, the CBC's flagship radio program, the Current, aired an interview with Kashmir Hill largely about the company, declining to detail Canadian experiments with

AI. Even with warnings from CCLA's Brenda McPhail that "we've got law enforcement agencies using it without confirming that the use of the tool is compliant with Canadian law. That's a big problem," the CBC ended by reporting that no police forces used the technology ("The End of Anonymity?," 2020).

It was not until the TPS set off a cascade of disclosures on 13 February 2020 that spokesperson Meaghan Gray emailed the press admitting that some of its members used the technology. Soon, police forces in Calgary, Halifax, Cornwall, and the federal RCMP began admitting they had used the tool (Tunney, 2019). While the Service de Police de la Ville de Montréal (SPVM) denied any application of Clearview AI, it stated that it might use similar third party softwares in the future. At the same time, Quebec's provincial police (Sûreté du Québec or SQ) confirmed its acquisition of facial recognition and fingerprint technology from the French company Idemia in August 2020, for \$4.4M (Ducas, 2020). The rush of disclosures prevented an early closure of the issue, instead bringing the American controversy up north.

The Clearview AI controversy thus presents the turning point for the governance of FRTs in Canada, demanding a strong response from policy makers and the setting for emerging policy consultations. Importantly, FRTs became present on the policy agenda largely due to an imported controversy, translated into a Canadian issue only after resistance by implicated actors. Facial recognition had been actively used in Canada dating back as far as 2014, but the Clearview AI scandal required governments and institutions to respond with the TPS, the OPC, and ETHI committee taking to the task through consultations.

Consultation on the OPC's proposals for ensuring appropriate regulation of artificial intelligence

The first of two policy consultations organized by the OPC on ensuring appropriate regulation of artificial intelligence took place between January 28 and March 13, 2020. While the consultation emerged out of an increasing attention by the OPC towards questions surrounding AI and biometric surveillance, the controversy around Clearview AI just days before the start of the consultations underlined their crucial timing.

With a specific focus on reforming the Personal Information Protection and Electronic Documents Act (PIPEDA), this online consultation asked for written responses to 11 policy proposals surrounding AI that the OPC’s analysts prepared in advance. In interviews, analysts appeared aware of the shortcomings of written submissions, which demand a lot of participants and risk, especially in the case of technology with such widespread application, missing key voices. The OPC did try to address these exclusions by adding two in-person consultations, one in Toronto and one in Montréal, the latter together with the CAI.

The OPC received 86 written submissions ranging from civil society, industry, academia, and the legal community (often legal actors tied to industry), which marked its largest consultation to date. The submissions are currently not publicly available due to privacy concerns, but an access to information request is underway. While the consultation received a large contribution from civil society actors, business interests and their legal representation disproportionately make up more than 50% of responses. The breakdown of participants offers a glimpse into how the online consultation acts as a composition, less a public, of actors able to participate by translating their positions into a format legible to the OPC.




Figure 1

Consultation on the OPC’s proposal for ensuring appropriate regulation of artificial intelligence

Based on the input of the submissions and the in-person consultations, the analysts prepared mini-briefing documents for each of the 11 proposals. Submissions were boiled down by actor to a simple yes or no regarding each of the proposals, with the commissioner having the final decision on which proposal to recommend to parliament. In addition, the OPC called upon Ignacio Cofone, assistant professor at the Faculty of Law, McGill University, to write an independent report reviewing the submitted responses.

In keeping with the original draft, the final regulatory proposal stressed a rights-based approach to AI and proposed reforms to Canada's privacy laws. Suggested reforms aimed at balancing Canadians' need for new and continued privacy rights without impending AI-driven economic growth. Overall, the regulatory proposal presented a sort of refinement of the original rights-based draft through the input of consultations, which sparks the question of whether consultations with predefined proposals foreclose any opportunity for radical departure from pre-envisioned policy positions.

OPC consultation on police use of facial recognition

The second of two policy consultations organized by the OPC focused on the police use of FRTs. It took place between June 10 and October 15, 2021. The consultation was sparked by the publication of the OPC's reports on the Joint investigation of Clearview AI, Inc. in February 2021, and the report on the investigation into the RCMP's use of Clearview AI in June 2021. The online consultation again called for written responses on draft guidance, which was developed by the OPC's policy and promotion sector in close collaboration with the investigative compliance department.

The OPC received 29 written submissions, which are currently not publicly available due to translation concerns. An access to information request is currently underway. In addition, the OPC held online roundtables with the Canadian and provincial human rights commissions and informal discussions with law enforcement agencies. The input of the former roundtable led to an additional roundtable that specifically involved affected and marginalized communities. These efforts were joined by the OPCs provincial counterparts with the Office of

the Information and Privacy Commissioner of Ontario holding stakeholder roundtables and the Office of the Information and Privacy Commissioner of British Columbia meeting with law enforcement agencies.




Figure 2

OPC consultation on the police use of facial recognition

Based on this input, the analysts summarized findings as a What We Heard report appended to the Recommended legal framework for police agencies’ use of facial recognition and a privacy guidance on facial recognition for police agencies in May 2022. Summarizing the consultation, one analyst recalls:

We invited organizations to read through the guidance and then meet with us, and then we structured it. So, we had a number of very specific questions that we asked them for their input on, and then heard what they had to say, and took notes and incorporated that into the revisions that we made to the draft (Anonymous, interview, 22.09.2022).

Similar to the previous consultation, the outcomes emphasized a rights-based approach, which already was apparent within the provided draft guidance. It underlines the OPC’s institutional

contingency as a fundamentally rights-focused institution. While the What We Heard report presents an added layer of transparency compared to the outcomes of the previous consultation, it is not clear which said what. This lacking transparency then extends the question of how the inner workings of consultations transform responses into the construction of issues and policy proposals.

Policy consultations by the Toronto Police Services Board

The Clearview AI controversy prompted the Toronto Police Services Board (TPSB) to initiate a public consultation on policy regarding the new (not current) use of artificial intelligence technology at the TPS in late 2021. The consultation was the result of a process of developing AI policy, which began in early 2021 with internal consultations at the TPS regarding the forces of FRTs and other AI tools. These internal consultations were jointly driven by Dubi Kanengisser, Senior Advisor, Strategic Analysis and Governance at the TPSB, and Deputy Chief Colin Stairs, the TPS's first ever Chief Information Officer.

Exchanges between Kanengisser and Stairs and internal consultations led to the development of a risk-based approach to FRTs and AI in policing. While the policy was developed independently, risk-based approaches to AI are commonly employed by law enforcement agencies, especially when dealing with the manifold AI systems that are already in use. An initial draft policy was informally passed to outside organizations with proven knowledge and leverage within the space of policing, including the Canadian Civil Liberties Association (CCLA), the Ontario Human Rights Commission (OHRC), the Law Commission of Ontario (LCO), and the Information and Privacy Commissioner of Ontario. The original drafts were met with lots of criticism by this initial round of feedback.

The draft policy was published on November 8, 2021, and opened for written responses via online consultations between November 15 and December 15, 2021. For a city with 2.7 M residents, the process received 49 applications with little response from civil society and business. Most notably the consultation gathered 31 submissions from individuals, the highest amount of public submissions in any of the consultations. In contrast to the OPC's

consultation, all submissions to the TPSB consultation are publicly available online and thus enable deeper insights into the nature of enrolled actors and emerging issues.

Figure 3

TPSB Use of new artificial intelligence technology policy - public consultation




Figure 3

TPSB Use of new artificial intelligence technology policy - public consultation

Four key issues that emerged in the submissions were: 1) the call for a moratorium and/or ban of FRTs until any policies are in effect, 2) the lacking independence of the TPSB to make this policy, 3) a broader definition of AI that encompasses all Automated Decision-Making Systems (ADMS), and 4) a lacking clarity of risk categories.




Figure 4

Sentiment analysis of submissions to the TPSB consultation on the use of new artificial intelligence technology policy

Even though the majority of individual comments raised issues with AI in the police services, the draft policy changed little in its final form. The new guidelines emphasized the TPSB's external partnerships more directly at the start and clarified procedures but did not, for example, ban certain applications, set a moratorium on AI systems currently in use, or define Extreme Risk Applications. Policy-makers at the TPS(B) questioned the quality of contributions as, at times, completely unrealistic when calling for an outright ban of AI. One interviewee stated that

the broader public one [consultation] is a lot less useful because it's just sort of Looney Tunes. People come out and say things like: 'we should never use AI ever, ever.' There are some good people who provide thoughtful feedback in that session, but often not what hasn't already been covered by the NGOs (Anonymous, interview, 18.08.2022).

Overall changes in the final policy did demonstrate some attention to comments, but neither responded to the substantial doubts raised about the technology nor did the TPSB elaborate the benefits of the policy when it was launched.

Following the adoption of the policy by the TPSB on February 28, 2022, the TPS moved towards its implementation via procedures which, rather than resulting out of the policy, co-evolved alongside the policy-making process for a quick implementation. The TPS posted these procedures for public consultation between November 24, 2022 and December 23, 2022. The procedures are being finalized at the time of writing.

Consultation (Expert Hearings) at the Standing Committee on Access to Information, Privacy and Ethics

AI was raised before the opposition-led Standing Committee on Access to Information, Privacy and Ethics (ETHI) as early as 2019, when the committee met to consider the Ethical Aspects of Artificial Intelligence and Algorithms. From April 30 to June 6, 2019, the committee invited contributions by prominent figures such as Yoshua Bengio, representing MILA, and Brent Mittelstadt of the Oxford Internet Institute, in addition to a policy brief submitted by Deloitte Canada. The Clearview AI controversy, growing concerns with racialized policing around the Black Lives Matter movement, civil liberties concerns with body-worn cameras and predictive policing, and recent investigations by the OPC all prompted the committee to initiate an in-depth investigation into issues raised by FRTs and AI.

The expert hearings took place from March 21 to June 16, 2022 and summoned 31 witnesses during nine sessions. In addition, the committee gathered eight policy briefs including four briefs of witnesses that also appeared during the hearings. As a parliamentary expert hearing the process differs quite starkly from the previous three case studies, focusing mostly on interviewing elites, experts, institutions, and civil society organizations. The breakdown of witnesses revealed the highest number of police forces, who often faced tough questions from the committee. While no members of the public participated, the presence of civil society actors as diverse as the CCLA, the International Civil Liberties Monitoring Group (ICLMG),

and the National Council of Canadian Muslims underlines a certain attention towards the representation of citizen concerns.




Figure 4

ETHI Hearings on the use and impact of facial recognition technology

The hearings addressed the issue of FRTs and AI in a broad manner, leaving a lot of ways for witnesses to engage with the topic. Compared to public consultations, these hearings are much more a political process with different political parties manifesting their own positions through the choice of topics that are being tabled, the witnesses that are being invited, and the questions that are being asked. The gravity of the application of FRTs by law enforcement is underlined by the fact that two of nine hearings were dedicated to questioning the TPS and the RCMP on their use of FRTs. The latter continued to remain evasive throughout both hearings, thus circumventing any form of accountability to the committee. As one interviewee recalls:

what we got was repeated tactics used by the witnesses both on the Government side and within the RCMP that deflected from any kind of real accountability (Matthew Greene, interview, 25.08.2022).

The committee's report on Facial Recognition Technology and the Growing Power of Artificial Intelligence was published on October 4, 2022 (Standing Committee on Access to Information, Privacy and Ethics, 2022). The report presents the results of the hearings with selected and detailed quotations from witnesses. It is organized around 19 recommendations for the Canadian government to take action on firmly underlining a rights-based approach to regulation. The report emphasizes the need "to reassure Canadians that the use of FRT or other AI tools in Canada is done responsibly and respect their rights." Most importantly, the report concludes with the general recommendation to impose a moratorium on the use of FRTs in Canada based on the feedback of expert witnesses.

On February 21, 2023 the Honourable François-Philippe Champagne, Minister of Innovation, Science and Industry, presented the government's formal response, specifically addressing how the upcoming Artificial Intelligence and Data Act (AIDA) as part of Bill C-27 addresses raised concerns.

Comparative breakdown of FRT consultations in Canada

In answering our first research question, the previous paragraphs and graphs show which actors are enrolled by policy consultations surrounding FRTs. Combined with our interviews these statistics also point towards how actors might be enrolled, which we discuss in the next section. More importantly, these statistics raise the question of which voices might have been absent or silenced by these processes, as well as why that might be the case. To understand this, in Summer 2022 we initiated an online survey targeted at civil society actors, most notably those representing the marginalized communities that are most affected by the deployment of FRTs in policing.

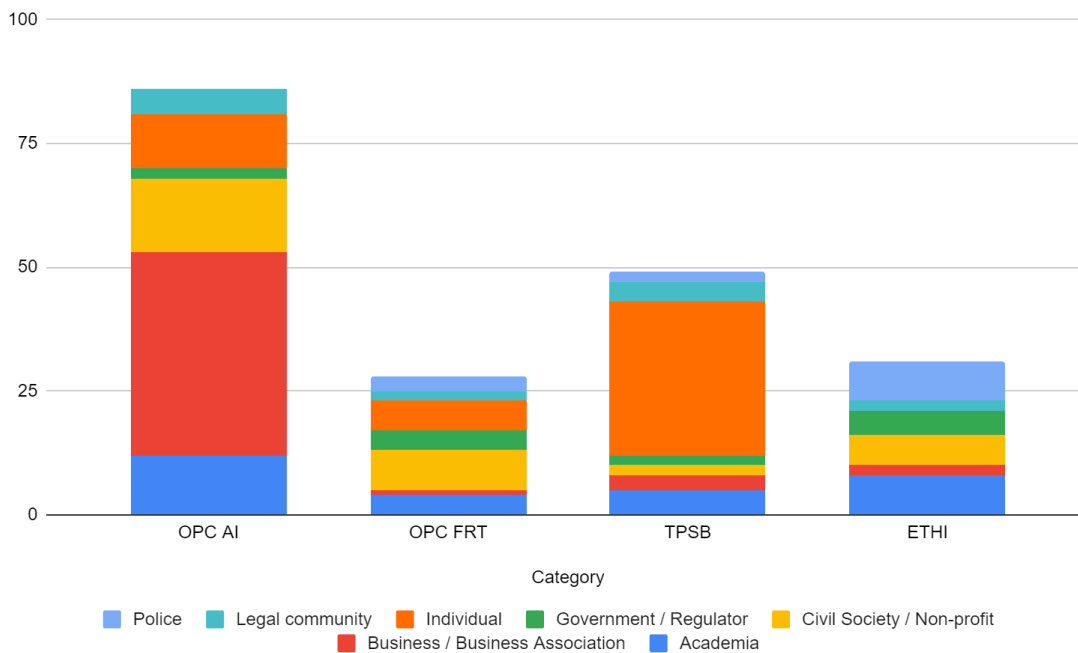


Figure 5

Comparative breakdown of consultations in the governance of facial recognition technologies in Canada

We received a total of five responses in English and French to our questions on how Facial Recognition Technologies (FRT) affect their organization and/or community; if FRTs are a concern in their locality; if and how FRTs should be regulated; the role of civil society in the regulation of FRTs; the role of consultations; and what might be missing in terms of regulation of FRTs by governments, companies, and/or civil society in Canada.

The results of our survey showed a split between liberties-oriented civil society organizations and organizations working directly with affected communities. For the former, FRTs are a major concern due to their (potential) impact on fundamental freedoms, privacy, freedom of association and assembly, and equality. For community organizations there is a concern for understanding how FRTs might affect their communities both positively and negatively. Responses showed an obvious knowledge gap between liberties-oriented and community organizations. The former responded extensively to our questions while the latter remained

brief and referred to their lacking expertise and resources, a common issue in participatory governance (Jones, 2022). In addition, liberties-oriented organizations referred repeatedly to their active involvement in addressing questions of FRTs and policing with several of them having contributed to policy consultations. Community organizations on the other hand neither contributed nor were invited to consultations, but rather emphasized a lacking public debate about these issues.

Analyzing the role of consultations in the governance of facial recognition technologies in Canada

In the following paragraphs we aim to answer the research questions two and three: 2) How do policy consultations enroll actors and construct issues surrounding facial recognition technologies? and 3) How do consultations affect concrete policy outcomes?

The four outlined cases underline how subjects (actors), objects (issues), and the modes of participation (political ontologies) are co-constructed through their consultative performances (Chilvers & Kearnes, 2020). Each consultation enrolled a different set of actors, addressed a different set of issues, and deployed different methods of engagement. We find that far from assembling the publics necessary for the democratic governance that these consultations outwardly promise, consultations instead perform the appearance of democratic governance.

Actors emerge through consultation and, for emergent issues like AI, depend on these participatory forms of governance to shape their opinions and positions (Pestre, 2008). The TPSB consultation, within its localized context, received a high percentage of written citizen submissions. Whereas the ETHI hearing had the power to summon witnesses selected by a political committee for testimony, the OPC put specific attention towards enrolling civil society actors by extending its call for written submissions towards roundtable discussions with civil society and community organizations. As one analyst recalls:

It's always going to be difficult to hear from everybody. Certainly for the online consultations, we tried to create lists of stakeholders that were very broad, and to reach groups that we knew likely would not provide us with written submissions, but perhaps would be open to attending a two-hour

meeting online and giving us their views on this piece of guidance and on a potential legal framework (Anonymous, interview, 22.09.2022).

Timeframes further influenced actors' enrollment. One of our interviewees mentioned that an ideal timeframe for consultation is between four to six months in order to allow for ample time for actors to develop solid contributions. All but one of the four cases reached the minimum of four months. The shortest time frame presented, the TPSB consultation, was a mere one-month, a feat that the TPS repeated despite wide-ranging critique about the limited and narrow format of a one month consultation. Seeing that resource and time constraints are the main hindrance to effective participation (Jones, 2022), the critique of policy makers at the TPS and TPSB that public contributions are lacking in quality might thus be considered a self-inflicted issue. As we stated elsewhere: "The TPS is the only force in Canada consulting on AI, which is commendable but such leadership comes with a responsibility to listen to feedback and develop effective ways to engage with the public" (Linder et al., 2023). If consultations shape publics, then the consultations left little time or energies to develop them.

Limited media coverage also undermined consultations. Newspapers and major media outlets did not cover these consultations with much depth. While controversies received wide media coverage, consultations did not. Why is this, especially if consultations offer the public a chance to engage rather than be political spectators? Pairing this with limited efforts by the government in publicizing consultations then accounts for an overall low number of submissions. As one interviewee referring to the TPSB consultation stated:

I spoke twice to TV channels, but nothing nearly sufficient to get public attention. [...] It is really unfortunate that there was a whole lot more media interest when we passed the policy than when we started this consultation. Probably you know it has a lot to do with just the general cynicism about these consultation processes. People think that either it's a done deal or it's just a stalling method. I'm hoping that, as we continue to have this process, more people in the media will see that this is actually a serious thing and we need their help and getting the word out (Dubi Kanengisser, interview, 13.07.2022).

The two consultations by the OPC and the consultation by the TPSB showed how strongly-preconditioned consultations are when posited around draft policies and guidance documents. In both cases, the draft policies posted for consultation presented a certain framing of the issues and proposed solutions at hand—a rights-based approach for the OPC and a risk-based approach for the TPSB. While the OPC’s rights-based approach can be traced down to its institutional legacy, the TPSB’s risk-based approach is commonly employed by law enforcement agencies. This shows how consultations are never spontaneous or a direct reaction to real-life events, but pre-planned with the issues carefully chosen and framed by the organizing actors (Thorpe & Gregory, 2010).

Limited consultations affected policy outcomes. Policy consultations diversified inputs but the final outcome was heavily preconditioned by existing policy positions (Jones et al., 2023).

While submissions are not publicly available for the OPC’s consultation, the What We Heard report published alongside its draft guidance for the police use of facial recognition are presented as a refinement to preexisting policy positions. In other words, The TPSB consultation on the other hand can be read as a resistance to public input. While the majority of submissions called for a moratorium and/or ban of FRTs, the resulting and adopted policy document resisted any mention or incorporation of these standpoints. Still, the consultation was viewed as a success, with one interviewee stating that

There was a lot of good stuff in there. And there was also more not really helpful stuff than I would have liked, but overall, I think it was a successful consultation process and I'm happy that we did it because the policy is better for it (Dubi Kanengisser, interview, 13.07.2022).

The question that lingers then is whether policy is better because of or despite public input. More importantly, one might question if the consultation was perceived to be successful because it did not necessitate major changes to the draft policy. In other words, holding a consultation in itself is viewed as a success, rather than how it addresses concrete policy outcomes.

Beyond the TPSB, calls for a moratorium and/or ban of FRTs in policing presented a widespread concern, equally emerging in submissions to the OPC and in front of the ETHI committee. In addition, a group of 29 civil society and community groups led by the International Civil Liberties Monitoring Group (ICLMG) publicly posted an open letter calling for a ban of FRTs as early as July 2020. Despite these calls, neither the OPC nor the TPSB included these in their final policy documents. Only the ETHI committee recommended the need for a moratorium of FRTs until adequate policy is in place. Both a call and an opportunity for banning the technology were possible at the outset of consultations. The window quickly closed. One analyst stated that

there's the kind of normalization, like resignation, to the idea that this is going to be used. I think, in early 2020 there was still a little bit of hesitancy around [...] should one accept outright that these technologies are going to be used by the police. [...] I don't know that it was taken as inevitable at that point early on that the police should be even allowed to use it. That was still kind of up for the debate. I think that for a lot of people that ship has probably sailed at this point. Where there's this kind of understanding it's no longer about yes or no, but like to the extent that they will what would be appropriate (Anonymous, interview, 22.09.2022).

While presented as central, it appears that consultations are merely one of a variety of factors that play into the governance process (Chilvers & Kearnes, 2020; Pestre, 2008). In the case of the TPS(B)'s ongoing investigations into the application of Clearview AI, questions of efficiency in delivering draft policies and procedures, the need to “appear to be doing something”, and not taking away powerful tools from the hands of police officers all played into the policy-making process. At the OPC the simultaneous investigations on Clearview AI, the RCMP and Cadillac Fairview, and the agenda setting by the privacy commissioner all played major roles in defining policy. In terms of the ETHI committee, interviewees described how the gathering was susceptible to partisan politics, as well as a stage for rising politicians to prove their grounds. Outcoming policy documents, such as the TPS(B) policy and procedures and the OPC's draft guidance, are exemplary of the foreclosure that preexisting policy

positions present in participatory processes and how consultations consequently swallow or sideline raised concerns (Jones et al., 2023).

In summary, our analysis underlines how public consultations serve as powerful closure mechanisms to controversies surrounding FRTs and AI. While presented as making “democracy more democratic” (Thorpe & Gregory, 2010), consultations are deployed in the wake of controversies and the anticipation of strong disagreements emerging on the issues at hand. While “controversies are situations where actors disagree (or better, agree on their disagreement)” (Venturini, 2010, p. 261), processes of public consultation in Canada appear to absorb these disagreements by creating a mechanism for contribution but not for debate. Common outputs, such as What We Heard reports, seem to be indifferent to highlighting disagreements, thus creating a politics without adversaries by manufacturing the illusion of consensus around dominant and preexisting policy goals (Elam & Bertilsson, 2003). In other words, public consultations weaken criticism by recycling it into providing the democratic justification of preexisting claims and policy positions. Participation thus presents a form of control and co-optation to both shape publics into accepting policy and shape policy to be accepted by publics (Thorpe & Gregory, 2010). Participation in the governance of FRTs in Canada is then a powerful example of the foreclosure of regulation by solution-first approaches to governance under the premises of Machine Learning Political Orders (Jones et al., 2023).

Conclusion

You have a government in particular [...] who prides itself on the language of consultation. You hear it in languages around reconciliation and anti-black racism. They love to consult. What are they doing with the consultation? There's a lot of arrogance within the halls of Parliament, of who knows best this kind of patronizing, [a] very elitist way in which we do politics in Canada that excludes and by design politically estranges people, so not only do they feel disempowered and left out, but completely disconnected from government (Matthew Greene, interview, 25.08.2022).

Canada had not one, but at least four consultations on FRTs and AI sparked by their controversial use in policing. Rather than building public trust, averting crises of expertise and democracy, and presenting more responsive and responsible governance mechanisms, consultations further complicate the personal relations that publics might have to AI and its governance. Technology policy consultation as tools that encourage deep public or civil society engagement might be intrinsically ineffective in their current state.

Participation needs to become more than depoliticized forms of consultation (Thorpe & Gregory, 2010). How different logics of public participation enroll publics, construct issues, and employ methods of participation directly impacts who might benefit and who might be adversely affected by AI (Chilvers & Kearnes, 2020). New, more equitable, diverse, and inclusive forms of public participation must be identified and implemented for the ethical and human-centered development of AI to create “ongoing relationships based on mutual benefit, reciprocity, equity and justice” (Sloane et al., 2020). Digitalization might well present the necessary tools to turn public participation from democratic legitimation and powerful closure towards public oversight of the machinery (Fourcade & Gordon, 2020b).

The status quo response to the consultations above belays deeper democratic risks caused by rapidly advancing AI. Spamming or astroturfing online submissions has been a persistent risk ever since three firms intentionally flooded a US Federal Communications Commission consultation with fake comments (New York State Attorney General, 2023). Extrapolating the spammed submissions, Nathan E. Sanders and Bruce Schneier (2023) warned in a New York Times opinion piece that, “Right now, the only thing stopping a ChatGPT-equipped lobbyist from executing something resembling a rhetorical drone warfare campaign is a lack of precision targeting. A.I. could provide techniques for that as well.” For a country that brags about being first to launch a national AI strategy, Canadian democracy seems to lag far behind when it comes to the next-generation of political bots hijacking these consultation processes (Dubois & McKelvey, 2019). Concerns for more robust forms of participation are emerging. To conclude with the words of one interviewee:

I think there's still room for improvement. The format of consultations, the way that we typically do consultations, tends to favor certain groups over

others. Asking people for written submissions, and that, you know, and asking them to answer ten very detailed and sometimes technical questions can be difficult for people. We're looking at how we can create different ways of soliciting input for different audiences, for different stakeholders (Anonymous, interview, 22.09.2022).

Chapter 5: AI and the Canadian Institute for Advanced Research by Robert Hunt and Théo Lepage-Richer (with research assistance from Nick Gertler)

Introduction

On March 22, 2017, the Canadian Institute for Advanced Research (CIFAR) issued a press release announcing that “the Government of Canada is funding a Pan-Canadian Artificial Intelligence Strategy for research and talent that will cement Canada’s position as a world leader in AI” (CIFAR, 2017a). A key component of the government’s broader effort to boost artificial intelligence (AI) research and development in Canada, the goal of the \$125 million strategy was, as the release put it, to “attract and retain top academic talent in Canada, increase the number of post-graduate trainees and researchers studying artificial intelligence, and promote collaboration between Canada’s main centres of expertise in Montréal, Toronto-Waterloo and Edmonton.”

The release humbly noted that “the program will be administered through CIFAR,” while less modestly claiming that “Canada’s global lead in AI is due in large part to the early support by CIFAR of a group of researchers from around the world, led by Geoff Hinton at the University of Toronto, for over a decade.” Those familiar with CIFAR’s early history might have raised an eyebrow at this capsule account. Known as CIAR⁴ at its launch in 1982, the institute’s inaugural research program was named Artificial Intelligence, Robotics, and Society (Brown, 2007). Research into the social and societal implications of AI was dropped before the program even began, but nevertheless, despite the press release’s decadal timeline, “artificial intelligence” had in fact been on the minds of the institute’s administrators and researchers from the start.

⁴ CIAR only became officially known as CIFAR in 2007, but for the sake of consistency we have chosen to use the CIFAR acronym throughout.

An institution offering a streamlined version of its own story (rather than the somewhat knottier history of AI research) in a press release accentuating its perspicacity is not terribly surprising. What is perhaps more surprising is the relative lack of controversy over the government earmarking \$125 million for the Pan-Canadian Artificial Intelligence Strategy (PCAIS) and tasking an organization that is neither a university nor a governmental agency with executing it. Apart from some mild misgivings about a Toronto-based organization wielding influence over the entire country's AI research ecosystem, media coverage of the announcement reported little consternation (Ball, 2017; Dandurand et al., 2022). Even the response from British Columbia—excluded from the strategy—was muted. How did CIFAR find itself in such a trusted position?

More broadly, how did a nonprofit organization become the central hub of a network connecting individual researchers, academic institutions and centres, governmental agencies, federal and provincial funding, multinational corporations, and start-up firms? This is not a story with unexpected twists but one characterized by ideological continuity and the establishment of institutional stability. Managing the PCAIS represents in many ways a culmination of CIFAR's founding goal of creating productive networks among academia, industry, and government. The story of CIFAR and AI in Canada consists of the settling of scientific controversy, the homogenizing of research agendas, and the centralizing and coordination of scientific research funding and industrial policy, both federal and provincial. CIFAR's history attests to a deliberate and successful intertwining of public and private interests and shows how scientific research, governmental funding, and the technology industry can become closely harmonized—though with the latter usually singing lead.

This chapter does not provide a comprehensive chronicle of CIFAR's forty-one-year history. Instead, we highlight significant moments and figures within the story of the institute becoming central to Canada's AI policy agenda. We begin by examining how CIFAR's approach to research networks was celebrated as a model for Canadian policy makers. We then pivot to analyzing how neural networks became the exclusive approach to AI at CIFAR. Next, we discuss how the institute came to embrace "artificial intelligence" as both a term and a specific set of techniques with seemingly infinite applicability in response to increasing

enthusiasm for machine learning within Silicon Valley. We then step back to show how, following the shuttering of its initial AI program in 1995, the appointment of Liberal Party policy maker Chaviva Hošek as president and CEO helped the institute achieve the financial stability needed to revive its AI research agenda. Throughout we call attention to CIFAR’s deep-seated commitment to bringing together academia, industry, and government—another kind of network that eventually led to the PCAIS.

1978	The idea for a Canadian Institute for Advanced Research is first introduced by John Leyerle, Professor of English at the University of Toronto
1979	A formal proposal is submitted to the University of Toronto
1980	A feasibility committee recommends that the institute be independent from the university and designed to connect researchers from universities across Canada
1982	The institute is formally established, with Fraser Mustard as its founding president (and first full-time employee)
1983	CIFAR’s Artificial Intelligence and Robotics (AIR) program is established, with William Tatton as its director
1984	CIFAR welcomes its first cohort of AIR fellows
1985	Zenon Pylyshyn succeeds Tatton as director of AIR
1986	Geoffrey Hinton joins AIR/CIFAR as a senior fellow and is provided a tenured appointment at the University of Toronto
1988	A first review of the AIR program concludes that the program should be renewed for another five years, but that greater emphasis on building links both between sites and with the industry should be put in place
1989	CIFAR inaugurates the Pre-Competitive Applied Research Network to support AI research in industry
1995	The AIR program ends
1996	Fraser is replaced by Stefan Dupré as CIFAR’s President and CEO
2001	Chaviva Hošek succeeds Dupré
2004	Hinton returns to CIFAR to launch and direct the NCAP program

2012	Alan Bernstein becomes President and CEO of CIFAR
2014	Yoshua Bengio and Yann LeCun take over direction of NCAP, now renamed LIMB
2017	The Pan-Canadian AI Strategy is announced; CIFAR moves to MaRS
2022	The PCAIS is renewed and expanded; Stephen Toope succeeds Bernstein; LeCun is replaced by Kording

Table 1: Significant Events in CIFAR History

CIFAR and the Network Model

CIFAR’s current reputation as the catalyst behind the rise of the machine model known as neural networks—a biologically inspired, statistical approach to machine intelligence—can be traced back to a series of “victories” won by some of its researchers from the early 2010s onward. In 2012, for instance, a team led by CIFAR affiliate Geoffrey Hinton disrupted the whole field of AI by winning the ImageNet Large Scale Visual Recognition Challenge⁵ with a neural network capable of classifying images at an error rate almost 50% lower than prevailing approaches to computer vision. Ten years later, be it in the popular (“How Canada’s Unique Research Culture Has Aided Artificial Intelligence,” 2017) or technical (Goodfellow et al., 2016) literature, CIFAR is now regularly credited with having fostered the current wave of AI research both in Canada and abroad. But CIFAR’s lesser-known, and arguably most palpable, contribution instead consists in having promoted another type of network—namely, research networks—as a privileged model for the organization of research in a country as diverse and decentralized as Canada.

While first conceived of in 1978 as the Canadian equivalent to the prestigious Institute for Advanced Study (IAS)⁶, CIFAR initially failed to secure the necessary endowment to reproduce the IAS model. After years of fruitless fundraising, the initiators of CIFAR developed a hybrid

⁵ ImageNet was a yearly image indexing competition evaluating algorithms for image classification, which took place between 2010 and 2017.

⁶ While symbolically associated with Princeton University, the Institute for Advanced Study is a fully independent research organization almost exclusively funded by private funding and philanthropic donors, whose mission consists in building capacity in fields with great technoscientific and industrial import.

model where research would be distributed across publicly funded universities yet organized around programs with strong appeal to both philanthropic and private donors.

After having surveyed the interests and priorities of key donors, the CIFAR scientific committee shortlisted three areas—astrophysics, population health, and artificial intelligence—as the most promising themes for the institute’s first research program. After almost a year of deliberation, which led to astrophysics being ruled out as already too mature a field for CIFAR to make an impact in and population health too broad, AI was selected as CIFAR’s inaugural program in 1982. Using a cross-affiliation model, the institute appointed rising figures in the field as fellows, who remained at their respective institutions while being required to closely collaborate with one another. Starting with a cohort of eight fellows—four at the University of British Columbia (UBC), two at University of Toronto (UT), and two at McGill University—CIFAR’s AI program set the stage for a decentralized, network-like approach to research management, with CIFAR at its centre.

While now widespread both in Canada and abroad, at the time this model was innovative from both a research and research-management perspective. As CIFAR’s president Fraser Mustard then noted, this network-like approach was seen as especially well adapted to tackle “some of the barriers” that inhibited “the full use of the country’s intellectual resources” (quoted in Brown, 2007, p. 68). For him, these barriers were “not only economic,” with regard to Canadian institutions’ limited financial resources compared to their American and European counterparts, but “also geographic” (quoted in Brown, 2007, p. 68), due to the sheer size of Canada’s land mass. In both cases, CIFAR’s unique model of creating new bridges, links, and connections across space, fields, and industries represented a solution to transcending economic limitations and physical distance alike.

So new was this model, and so widely recognized were these barriers by the Canadian political class, that in 1982, when the Royal Commission on the Economic Union and Development Prospects for Canada—also known as the Macdonald commission—was appointed, CIFAR provided one of the main case studies evaluated by the commissioners. Convened during Pierre Elliott Trudeau’s premiership, the Macdonald commission was charged with identifying ways to ensure Canada’s long-term economic prosperity in the face of both social and

technological change. Beyond its promotion of free trade, for which it is now famous, the commission advocated for the federal government's renewed involvement in fostering technological innovation and socioeconomic equality. In fact, it predicated the government's very capacity to respond "to ever more encompassing notions of equality" on the development of "new instruments and administrative technologies" capable of producing the wealth necessary to fund new welfare programs.

Of all the technologies it surveyed, the Macdonald commission identified the computer as uniquely suited to achieve these socioeconomic aims. Singling out AI as one of its most promising applications, the commission promoted its capacity to "substitute . . . human mental and physical efforts" (Macdonald, 1985a, p. 117) as an opportunity to free Canadians from physically straining jobs in industries such as manufacturing and resource extraction. To identify how best to catalyze the development of AI, the commission selected as one of its experts CIFAR's Fraser Mustard. Emphasizing how "nations with strength in one sector of a rapidly developing field have entrée to all its sectors" (quoted in Macdonald, 1985a, p. 117), Mustard recast AI as a national imperative that should be pursued to ensure future prosperity. To do so, he recommended that the government increase "the overall fund of knowledge and techniques open to their industries and institutions" (Macdonald, 1985a, p. 117) with, however, little emphasis on how that funding should be allocated.

But beyond Mustard's testimony, it is the model embodied by CIFAR that appealed to the commissioners. Clarence Barber, the commission's vice-chair, was one of CIFAR's early initiators, having co-developed the network-like model that came to characterize the institute. Described in the commission's final report as a "private sector initiative undertaken with public sector support," CIFAR represented a promising new model to "help Canada move to the forefront of technological innovation" (Macdonald, 1985b, p. 205).

More specifically, the institute's approach to network-building attracted the commissioners, who saw in it the blueprint for a new approach to not only industrial policy but also nation building. Drawing from CIFAR's model, the commission recommended the establishment of "several technology-development centres . . . in computer technology and related fields" (Macdonald, 1985b, p. 145) to not only develop a coast-to-coast pool of expertise, but also

create closer links across provinces, in a way that would transcend their individual interests and priorities. In continuity with the New Public Management movement and other more “business-like” models to governance in vogue at the time, the commission put forward a highly decentralized approach to both technological and national development—one where the fostering of new and more flexible links across universities, companies, and local governments would allow Canada to achieve key economic and political goals nationally.

Apart from the pursuit of a free trade agreement with the United States, which remains the commission’s defining recommendation, the creation of the Networks of Centres of Excellence (NCE) program was arguably the second-most important policy outcome. Created under the administration of Brian Mulroney with an approved budget of \$240 million, the program aimed to create networks of partnerships across provinces around research topics that were deemed strategic to Canada’s long-term economic prosperity. While distinct from the NCEs, CIFAR provided a clear—and acknowledged—inspiration for the program (Industry, Science and Technology Canada, 1989) and became, from its creation onward, a regular beneficiary of federal grants. It is thanks to that new, steady source of funding, which freed the CIFAR leadership from the near constant struggle of securing financial support, that CIFAR expanded its mandate by not only setting upon new research programs, but also bringing in new contributors—as well as new perspectives—into its AI program.

Fraser Mustard	Founding President	1982–1996
J. Stefan Dupré	President and CEO	1996–2000
Chaviva Hošek	President and CEO	2001–2012
Alan Bernstein	President and CEO	2012–2022
Stephen Toope	President and CEO	2022 to present

Table 2: Presidents of CIFAR

AI, Hinton, and the CIFAR Model

This transition was spearheaded by the appointment of cognitive scientist Zenon Pylyshyn as the new director of AI research. While his predecessor William Tatton, a neurologist from

Toronto Western Hospital, had mostly selected fellows from a similar background, Pylyshyn focused on bringing in representatives from all the main schools of thought and approaches to AI research. This policy contrasted with the framework used in other national settings to allocate research funds in the field of AI. As documented by Jon Guice (1998), for instance, government institutions in the United States each invested in the strand of research that seemed best aligned with their own needs and priorities, if only to justify their research budgets remaining distinct from one another. But, as he noted, Pylyshyn believed that Canada could only compete with other nations in that sphere by fostering greater “interaction across disciplinary, institutional and regional boundaries” (quoted in Brown, 2007, pp. 63–64) and thus adopted a research policy that promoted collaboration in a way that, at times, went against his own disciplinary commitments.

Throughout his tenure, Pylyshyn expanded CIFAR’s AI program to welcome new fellows, such as control theorists Peters Caines and George Zames and symbolists like John Mylopoulos and Hector Levesque, and to secure appointments at Canadian universities for scholars from underrepresented fields. These included Wolfgang Bibel, a German logician who joined the UBC faculty, and the American roboticist John Hollerbach, who was provided with an appointment at McGill. The resulting network, which spanned sites and approaches, built on the assumption that disagreements, controversies, and opposing agendas could ultimately be resolved if their representatives were made to inhabit the same network.

Pylyshyn’s commitment to turning CIFAR’s AI program into a space where distinct approaches could at times collaborate, and at others compete, was most notably illustrated by the recruitment of one of his own intellectual antagonists, Geoffrey Hinton, then at Carnegie Mellon University. While Pylyshyn’s work posited that the type of computing relevant to the study of cognition and, by extension, machine cognition necessarily involves “operations on symbols” (Pylyshyn, 1980, 1986), Hinton argued that the processes and functions underpinning cognition were reducible to both the material organization of the brain and the mathematical principles it embodies.

At its core, Hinton’s work aimed to explain how incredibly complex tasks like visual perception, speech, and spatial orientation could be computed by physically finite systems like

the brain. Focusing primarily on the model known as “neural networks”—a biologically inspired approach that aimed to reproduce in either hardware or software the statistical principles embodied by neurons—Hinton worked on developing highly parallelized computing architecture inspired by the structure of the brain, which could perform tasks exponentially more complex than the structure computing them (Fahlman & Hinton, 1987; McClelland et al., 1986). Hinton’s appointment in 1986 as a CIFAR fellow at the University of Toronto coincided with the publication of Pylyshyn’s well-known critique of the connectionist movement, of whom Hinton was one of the main representatives, in a piece titled “Connectionism and Cognitive Architecture.” In it, Pylyshyn rejected Hinton’s neural networks as attractive only to those who approached cognition from the perspective of its computational implementation and thus assumed that cognition can “only be understood if we study it as neuroscience” (Fodor & Pylyshyn, 1988, p. 4).

Despite their disagreements, Pylyshyn endorsed the recruitment of Hinton, who was provided with a tenured appointment at UT. Beyond his commitment to bringing a wide range of disciplinary perspectives together, Hinton’s recruitment was part of Pylyshyn’s efforts “to facilitate the transfer of knowledge and the development of a strong applied research sector” (Brown, 2007, p. 64). For in the years preceding his appointment, Hinton had been unusually successful at securing financial support from both governmental and private bodies to pursue applied research. This feat was especially notable as, except for a small network of scholars affiliated with the Massachusetts Institute of Technology, most actors involved in AI research in the United States were struggling to find funding from both the public and private sectors.

In fact, following Hinton’s appointment, CIFAR’s proximity with, and reliance on, both industrial and public sponsors, which epitomized the type of intersectoral ramifications underpinning both its network approach and its funding model, translated into a gradual homogenization of CIFAR’s research agenda. While the first five years of CIFAR’s AI program were characterized by the creation of new links among neuroscience, computer science, and other fields to develop a more comprehensive theory of intelligence (1982–1987), the subsequent era was defined by the gradual adoption of industry and/or government partnerships as the main source of validation for the type of AI research carried on by CIFAR.

For as formulated in a contemporaneous landmark report from the Canadian Ministry of State for Science and Technology, the optimization of both industrial and government operations was at the time perceived as AI's most promising application (MacKinnon, 1987, p. 11). Upon his arrival in Canada, Hinton directly benefitted from this perception by securing several grants and lucrative partnerships with both provincial and federal bodies. Compared to his previous work—which, funded primarily by the US military, had focused on abstract explorations of computing functions in highly parallelized systems—his CIFAR-sponsored research instead centered on the automation of processes with high operational import.

With grants from the Information Technology Research Center to research speech recognition, for instance, Hinton developed a time-delay neural network architecture that, while optimized to identify phonemes in sound samples of different lengths, was more broadly designed to process any “imprecisely prepared training examples” (Lang et al., 1990, p. 42). For the Canadian Institute for Robotics and Intelligent Systems, which financed work on interactive interfaces, Hinton applied neural networks to the development of a gesture-to-speech interpreter (Fels & Hinton, 1993, p. 7) with the broader aim of improving neural nets’ “non-linear mapping” of complex data (Hinton & Frey, 1995, pp. 1–4). With funding from the Natural Science and Engineering Research Council, Hinton developed a neural network simulator, whose representation of neural nets’ learnings aimed to facilitate their application to abstract tasks. In all these cases, Hinton’s research was characterized by a balance between highly targeted applications and the development of increasingly generalized learning architectures, to the point where his models’ validity came to be attached to their execution of minute, operationally relevant tasks.

From CIFAR’s perspective, these sponsored projects constituted a major success, as they directly embodied the collaborative, decentralized, and, more importantly, self-organizing principles of its network approach. As a review committee contemporaneously noted, CIFAR’s interactions “with industry and the transfer of knowledge to it . . . [had remained] minimal” (Brown, 2007, p. 66), so Hinton’s work stood out as a welcome exception. In the following years, the CIFAR leadership increasingly focused on scaling its efforts in knowledge translation and, in the name of “realizing its goal of co-operative research between Fellows . . . at different

nodes” (Brown, 2007, p. 66), reorganized its AI program specifically around the type of work pursued by Hinton.

On the one hand, CIFAR appointed several new fellows whose work not only focused on neural networks but was also supported by industry and government partnerships like those developed by Hinton. These included scholars such as UBC’s David Lowe, who used similar models as Hinton’s to perform computer vision tasks, and UT’s John Tsotsos, who also worked on computer vision. This growing emphasis on partnership building and applied research also translated into the inauguration of a new type of appointment: the “AIR industrial fellows” given to researchers working in industry as well as the creation of a new research network around industrial applications of AI. With funds from dozens of industry partners and, later, federal grants, CIFAR established the Pre-Competitive Applied Research Network, or Precarn, to support industry-based AI research in fields such as manufacturing, energy, and mining.

On the other hand, scholars from other branches of AI research that were less attractive to government and industry saw their CIFAR sponsorships dry up, sometimes after more than a decade of support. In 1995, for instance, Hector Levesque’s appointment as the AI program’s co-lead at UT was terminated, leaving Hinton as the sole overseer of CIFAR’s work in Toronto. Similarly, fellows like Wolfgang Bibel, who had been painstakingly—and at great cost—recruited from abroad, found themselves increasingly isolated at CIFAR and ultimately left Canada altogether.

While initially conceived of as a decentralized, spatially dispersed, and disciplinarily diverse model of research management, CIFAR’s network approach ultimately evolved into a mostly homogenous model where links with, and to, both industry and the public sector overshadowed interdisciplinary ones. Not only that, but the institute’s initial ambition to do away with political considerations by building new bridges across both linguistic divides and provincial bodies gradually gave way to the usual state of things, with Canada’s largest research universities—namely, UT, McGill, and UBC—receiving the lion’s share of CIFAR funding and support.

In the image of neural networks themselves, this network approach thus aimed to circumvent the limited resources at CIFAR’s disposal by increasing the number of links and connections

across all the sites—or nodes—involved. And like neural networks, this approach gave greater weight to certain nodes and, while distributed, reproduced a highly hierarchical model akin to the one it initially aimed to replace. Instead of a decentralized model of research management, CIFAR’s network approach thus exemplified how diversity and difference can be gradually reduced in distributed systems, with homogenization unfolding at the same time as new connections and links are established.

AI Research at CIFAR in the Twenty-First Century

Geoffrey Hinton would be the key figure around whom CIFAR’s support for research into AI would initially revolve in the twenty-first century. But CIFAR’s involvement in AI came to a dramatic, though ultimately temporary, halt in 1995, when the institute eliminated the Artificial Intelligence and Robotics (AIR) program due to recurring challenges in securing funding, particularly from the federal government, and the perception that it had become redundant with its offshoot Precarn (Brown, 2007, p. 205). These financial crises also prompted the resignation of Fraser Mustard, who was replaced by Stefan Dupré. Efforts to return AI to CIFAR’s portfolio of programs eventually succeeded, under the leadership of Dupré’s successor Chaviva Hošek, nine years after the elimination of AIR.

William Tatton	Director of Artificial Intelligence and Robotics	1983–1985
Zenon Pylyshyn	Director of Artificial Intelligence and Robotics	1985–1995
Geoffrey Hinton	Director of Neural Computation and Adaptive Perception	2004–2014
Yoshua Bengio	Co-Director of Learning in Machines and Brains	2014 to present

Yann LeCun	Co-Director of Learning in Machines and Brains	2014–2022
Konrad Kording	Co-Director of Learning in Machines and Brains	2022 to present

Table 3: Directors of CIFAR’s AI Programs

Hinton returned to CIFAR to launch the Neural Computation and Adaptive Perception (NCAP) program in 2004, marking a momentous year in the history of the institute’s involvement in AI development in Canada—though it perhaps did not appear quite so epochal at the time. While reflective of the disciplinary homogenization discussed in the previous section as well as the imprint of Hinton’s personal interest in the brain, CIFAR’s early descriptions of the program were nonetheless modest in recounting its aims, exhibiting equal enthusiasm about NCAP’s capacity to shed light on human *and* artificial intelligence, even while they eschewed the latter term. According to an archived CIFAR webpage, NCAP’s “two main goals” were:

to make major advances in understanding how the brain processes information (or learns) and to find out how to translate that understanding into the development of better artificial visual systems. As the second goal suggests, the program is approaching these challenges through the study of vision, and more specifically, the ability to recognize and make distinctions among three-dimensional objects. (*Neural Computation & Adaptive Perception Program*, 2005)

The order of these goals would eventually be reversed in the new name NCAP would adopt in 2015: Learning in Machines and Brains (LIMB). Not solely a shift in nomenclature, the new name also followed a change in leadership in 2014, when Yoshua Bengio and Yann LeCun took over as codirectors of the program. Their more ambitious and grandiose title attested to the trajectory of CIFAR’s embrace of “AI” in the post-AIR era from initially tentative to evangelical—as we will see, by the 2020s, AI was at the heart of CIFAR’s public image.

But they began much more reticently. CIFAR’s annual reports following the launch of NCAP make no mention of the term “artificial intelligence” until 2014, and even that first allusion refers to Big Tech’s interest in AI rather than an explicit assertion that CIFAR-affiliated researchers study or make AI. Deep learning and neural networks remained the preferred terms—an accurate reflection of CIFAR’s contracted interest in exclusively developing neural networks and applying machine learning to specific tasks. Indeed, for nine years CIFAR’s reports circumspectly avoid connecting the program to any broader discourse around artificial intelligence. And yet, in 2017, only three years after that initial demure reference, CIFAR would be tasked with managing the Pan-Canadian AI Strategy.

While their embrace of the term “AI” may have been hesitant, CIFAR was not modest about deep learning’s potential or its appeal to industry. The span of time between the launch of NCAP and the inauguration of CIFAR 2.0 in 2012 was marked by a continual expansion in scope of the applications of deep learning from recognizing three-dimensional objects to identifying and classifying human gaits, speech, and emotions to genome analysis and beyond. Reflecting the program’s intellectual roots in neurology, NCAP’s original focus was on computer vision. The program’s purported interest in the workings of human brains has endured (e.g., it has been maintained in the recent name change), though the extent to which NCAP or LIMB has contributed to neurological research is hard to determine based on CIFAR’s own reporting. While commercial applications and industrial take-up of computer vision are regularly touted in their annual reports, the utility of NCAP’s insights to neurologists is not clear.

Having made a significant leap in object recognition, NCAP researchers began to look for other kinds of objects for their machines to gaze upon. In 2009, the institute’s annual report highlights recent NCAP research on gait analysis, replication, and prediction, which they foresaw as useful to “animation companies” and in “video surveillance” (CIFAR, 2009, p. 17). With the leap to prediction the conceptual link between human and computer vision has already become strained. Rather than using computers to better understand humans, the goal seems to have silently shifted to developing machines that surpass them. And with the casually apolitical mention of “surveillance”—the report does not ask who is being watched or by

whom—CIFAR’s congenital preference for capitalizing on AI research over investigating its effects on society is evident.

The following year, NCAP added recognizing speech to its list of objectives (CIFAR, 2010). Once again, corporate interest in the research is foregrounded: “Speech recognition groups at IBM and Microsoft are now actively collaborating with NCAP researchers to make better speech recognizers by using ‘deep’ learning” (CIFAR, 2010, p. 6). That theme continued two years later, when their annual report noted:

Over the last year, the deep neural networks developed . . . by members of the NCAP program have started to replace traditional models used to decipher speech, as leading speech research groups at Microsoft, Google and IBM have shown that these deep neural networks are significantly more accurate. For example, the new Android 4.1 software uses a deep neural net as its acoustic model and Microsoft has also deployed this new approach for voice search. (CIFAR, 2012)

Research into the controversial field of emotion recognition—a technology whose shaky empirical and conceptual foundations have been critiqued from a variety of perspectives (Barrett et al., 2019; Hunt, 2018; McStay, 2018; Stark & Hutson, 2021)—soon followed (CIFAR, 2013, p. 13). Further applications for deep learning (e.g., genome analysis) only grew from there. Once the PCAIS had been announced in 2017—which put AI at the very heart of the institute’s mission and identity—AI applications began featuring in other programs in the natural and social sciences, including Brain, Mind, and Consciousness; Bio-Inspired Solar Energy; Genetic Networks; and Institutions, Organizations, and Growth (CIFAR, 2017b). Across nearly four decades, CIFAR’s inaugural interest in AI research had narrowed and homogenized into a restrictive focus on neural networks with the potential for industry-friendly applications. By 2017 this interest had metastasized, spreading across the institute’s programs, with deep learning treated as an all-purpose methodological panacea. In the institute’s 2018–19 annual report, NCAP was hailed as CIFAR’s most exemplary research program and was credited with having “unleashed the current worldwide revolution in AI” (CIFAR, 2019b).

Along with this widening embrace of AI as a crucial all-purpose research tool, the PCAIS also served to revive CIFAR’s foundational interest in the societal effects of AI, resulting in programs such as AI and Society (<https://cifar.ca/ai/ai-and-society/>) and a series of policy workshops. Gauging the seriousness of these efforts is tricky. The difficulty of simultaneously promoting the industrial adoption of AI and claiming it as force for the improvement of society while demonstrating concern for its impact makes the institute’s politics incoherent. For example, in their 2018–19 annual report, the “main goals” of the PCAIS are articulated as: “to maintain Canada’s position as one of the world’s leading nations in AI science and its applications in the economy, in government and in other areas of science; and to address the societal impacts of this disruptive new science” (CIFAR, 2019b, p. 29). An almost comic lack of balance between these goals is revealed on the following page, where we learn that CIFAR AI Chairs are

pioneering new research in fields such as: AI for drug discovery, AI for genomics and disease prediction, AI for materials discovery, artificial neural networks, autonomous vehicles, computer vision, deep learning, human-AI interaction, natural language processing, precision medicine, reinforcement learning, security and privacy, [and] societal implications of AI. (p. 30)

Alphabetical order may be the reason “societal implications” are consigned to the end of the list, but much of CIFAR’s interest in the effects of AI on society amounts to touting its societal *benefits*—naturally, an industry-friendly approach—and attending to its real or potential harms comes across as more of an afterthought. Likewise, while the PCAIS funding research into AI’s impact on society is undeniably an improvement over the institute’s earlier failure to maintain a social component in AIR, the differing levels of investment are telling (see Table 4, where the most 2023 figures indicate that the chairs program received nearly forty times the amount of funding as AI and Society).

Fiscal Year	AI Chairs	AI & Society
2019	\$14,143,000	\$368,000

2020	\$11,865,000	\$424,000
2021*	\$9,896,000	\$40,000
2022	\$15,431,000	\$256,000
2023	\$17,075,000	\$436,000

*Due to a change in CIFAR’s fiscal year, FY2021 only contains nine months.

Table 4: Comparative CIFAR Spending

The Path to PCAIS

As detailed in the previous section, CIFAR’s longstanding support for research into neural networks helps to partially explain the Trudeau government’s decision to task them with managing the PCAIS. Nonetheless, moving from a lack of federal funding leading to the elimination of AIR in 1995 to CIFAR being granted responsibility for a \$125-million national AI research program is a remarkable comeback. It is not possible to chronicle the full organizational history of CIFAR here, but a few significant events help to illuminate the changing relations between the institute and the federal government that would culminate in the PCAIS.

Stefan Dupré began attempting to resurrect AI research at CIFAR almost immediately upon taking on the presidency. The first attempt—focused on “collaborative systems” and led by UBC professor Alan Mackworth—failed to pass the institute’s Research Council, with some members finding its mission vague and, notably, questioning “why CIAR should pursue this proposal when the very strong groups in artificial intelligence at UBC, U of T, and McGill, which the earlier program [AIR] had created, could do this kind of work” (Brown, 2007, p. 225). While the goal of starting a successor program to AIR endured, obtaining enough funding to back a new program continued to be a challenge. Despite successfully developing the institute’s endowment, Dupré retired before the objective of winning dependable, long-term government support was met.

A notable shift in the institute's relationship with the government occurred in 2001, when Chaviva Hošek took over as president. Like Dupré and founding president Fraser Mustard, Hošek had prior experience in academia, having worked as a literature professor at the University of Toronto for thirteen years (CIFAR, n.d.). But unlike her predecessors, Hošek had left academia and moved decisively into politics many years before leading CIFAR. Standing for the Liberal party, she was elected to Ontario's provincial legislature in 1987 before being defeated in 1990 (Calleja, 2001). From there, she joined the federal Liberal party apparatus, first becoming head of their Research Bureau (Jeffrey, 2010) and then working in Prime Minister Jean Chrétien's office as director of policy and research. While holding those positions, she was credited with writing or cowriting the party's platforms in 1993, 1997, and 2000.

Described as a "former feminist lobbyist and policy wonk with a PhD in English literature" in a Canadian Business article about her becoming CIFAR president (Calleja, 2001), Hošek ostensibly made for a somewhat surprising choice for the role. Research in the humanities has never been anywhere near the top of CIFAR's agenda. Indeed, the profile makes much of the incongruity of a scholar adept at discussing deconstructionism leading a "scientific research network." But as Dawn Calleja, the profile's author, puts it, "[CIFAR] needed someone with connections to the country's elite," and Hošek's ostensible lack of direct experience with scientific research mattered less than the "huge web of contacts in government and business" she had developed during her political career. The institute's relatively low public profile and the intangibility of its decentralized, network-fostering model made fundraising a challenge. Hošek's mandate was to communicate the institute's strengths to potential industrial and governmental partners.

Despite the ostensible exoticism of her literary background, Hošek was a logical fit. Working in the PMO, she had been a prominent advocate for developing federal policy for and federal investment in scientific research (Axelrod et al., 2011). During her tenure, the federal government launched the Canada Foundation for Innovation and the Canada Research Chairs program. Born from the need to spend a budget surplus in a politically palatable manner, these two initiatives reflected a newfound willingness on the Canadian government's part to

intervene in the administration of universities, which fall under provincial jurisdiction, and a desire for those interventions to have clear benefits for the Canadian economy. Much like CIFAR, both programs reflected fears about exceptional Canadian researchers leaving for greener pastures abroad as well as a broader anxiety about Canada falling behind in a global race for innovation.

Evidence of CIFAR's improved standing with the federal government arrived in 2002, when it was announced that:

Ottawa would give CIAR a grant of \$25 million over five years, with a provision that the institute raise a matching \$25 million during the same period. The implications of the new agreement were enormous: the grant signalled strong recognition by the federal government of the importance of CIAR's research programs to Canada. The new grant would also have a major positive impact on private sector fundraising and transform the framework of financial stability built by Dupré into a firm foundation for the future. (Brown, 2007, p. 261)

While critical to the immediate financial health of the institute, this sizeable federal grant both laid an economic foundation and provided a strategic blueprint for the decades ahead. Going forward, a mix of renewing federal and provincial grants combined with private donations and investment income would essentially guarantee CIFAR's ongoing fiscal and institutional stability—and it would cement the institute's commitment to bringing together and mobilizing public funds, private donations, and scientific research. This financial and operational model would culminate in the PCAIS.

CIFAR 2.0, the PCAIS, and the AI Boom

In their public discourse, CIFAR's dedication to public-private partnerships takes on the aura of a high-minded philosophical principle. As they put it in their 2013 annual report:

“Partnerships have been part of the CIFAR model since it was created in 1982. The bold vision of CIFAR 2.0 now puts that concept at the very core of its operations. Increased partnering

affords CIFAR the opportunity to realize all aspects of its ambitious mission” (p. 39). CIFAR 2.0 was new president Alan Bernstein’s strategic plan to broaden the institute’s global reach. Defined by six goals, the strategy primarily consisted of building collaborations with like-minded domestic and international organizations (e.g., Banff, SSHRC, and France’s Inria) and the simultaneous recruitment of international scholars to and retention of Canadians within its research programs, another of the institute’s founding missions. As such, the “2.0” designation somewhat exaggerated the novelty of this strategy, which broadly speaking merely reemphasized certain aspects of the institute’s original mandate. Indeed, as the 2013 report also noted, “CIFAR has always been a partnership between the public and private sectors” (p. 44).

But in certain significant ways, CIFAR in the 2010s was a changed institution. As discussed previously, its embrace of deep learning across a widening number of programs represented a reversal of its original commitment to open-ended inquiry, pure research, scientific diversity, and interdisciplinarity. The years immediately preceding the PCAIS saw some significant changes in CIFAR’s AI program. As previously mentioned, Geoffrey Hinton stepped down as NCAP’s director in 2014, and Yoshua Bengio and Yann LeCun took over as the program’s codirectors. That same year, Hinton also began working for Google and LeCun for Facebook.

While public and private partnerships were always part of CIFAR’s modus operandi, a significant scaling-up took place in the 2.0 era that matched the institute’s improved financial health. In the 2014–15 fiscal year, for example, the institute received more private donations than in any of the previous seventeen years (CIFAR, 2015), and Bernstein’s promise to forge new partnerships was kept through collaborations with various governmental and industrial partners, including a LIMB summer school program involving both universities and corporations (e.g., Samsung, Google, and Amazon). In 2016, CIFAR’s annual report boasted that “revenue from all sources grew to \$19 million, an increase of 27 per cent over the previous year” (CIFAR, 2016, p. 16); and in 2017, the PCAIS was announced with an attendant, significant increase in federal funding. Driven by the enormous success of deep learning in the technology industry, CIFAR’s boom years had begun in earnest.

During this era, the institute’s positioning of itself as a field for conjoining research, economic, and policy interests became totally explicit. For example, in 2016 they bragged that CIFAR had

“received the inaugural Ideas Award from the Creative Destruction Lab at the Rotman School of Management in Toronto. The award recognized CIFAR for having a first-order impact on Canada’s competitiveness in the world by advancing new ideas in science and technology” (CIFAR, 2016, p. 9). In the wake of the PCAIS, the institute’s leaders would define its role through statements like this from PCAIS Executive Director Elissa Strome:

The success of Canada’s national AI strategy rests on two fundamental components: people and partnerships. This report showcases some of the amazing researchers and leaders from across our AI ecosystem, who believe deeply in the opportunity that AI has to make a real difference in people’s lives around the world. The only way that we will fully realize that opportunity is by working together, and that is why partnerships between government, academia and industry, as demonstrated across the CIFAR Pan-Canadian AI Strategy, will be the key to our success. (CIFAR, 2019a, p. 3)

Strome was well situated to frame CIFAR’s mission in these terms. Hired by CIFAR in 2017, she had previously worked as the Director of Strategic Initiatives at the University of Toronto’s Office of the Vice President, Research and Innovation, and as the Executive Director of SOSCIP, a government-funded research and development consortium whose “mission is to bring together industrial partners and academic researchers and provide them with sophisticated advanced computing technologies and expertise” (SOSCIP, n.d.). An acknowledgment in CIFAR’s 2020 report on the PCAIS cuts to the heart of the matter with nearly comic bluntness: “CIFAR’s leadership of the Pan-Canadian AI Strategy is funded by the Government of Canada, Facebook and the RBC Foundation” (CIFAR, 2020a, p. 31). The prospect of these various “partners” holding conflicting approaches, goals, or methods is never broached, nor are the differences in power between the state, multinational corporations, and individual researchers.

Reflecting their newly public and prominent role as leaders of the PCAIS—making them both recipients of relatively lavish federal funding and synonymous with the emerging, uncertain, and potentially dangerous technology of AI—CIFAR began to present “trust” as central to their brand, contending that “by building long-term, interdisciplinary, global communities of

collaboration, CIFAR provides the world's top researchers with an unparalleled environment of trust, transparency and knowledge sharing" (CIFAR, 2019b, p. 1). Outgoing chair of the board of directors Barbara Stymiest expanded on this theme, broadening it to include the institute's partners and writing that "at the core of these collaborations is trust, which CIFAR has built with leaders, supporters, and researchers from around the world" (p. 2).

Beyond these loftier proclamations, the PCAIS fundamentally transformed CIFAR as an institution. In their public discourse, the strategy is presented as one of the institute's "five interconnected pillars of excellence," which means it enjoys equal billing with all of their research programs combined (CIFAR, 2019b, p. 4). Perhaps most importantly, the PCAIS's financial effects were momentous: in 2020, the PCAIS accounted for 53% of CIFAR's revenue and 54% of its expenses. And in June 2022, the strategy was renewed and expanded: "In the 2021 federal budget, the Government of Canada affirmed its confidence in CIFAR's leadership of the Pan-Canadian AI Strategy by announcing more than \$443 million over 10 years for the second phase—including \$208 million to CIFAR" (CIFAR, 2022, p. 7).

CIFAR had always carefully balanced a form of soft, protective nationalism with a global outlook. But having increasingly become an auxiliary arm of federal industrial policy, CIFAR now publishes jingoistic statements such as: "Canada is one of the great AI nations, known for our pioneering excellence in the fundamental science of deep learning and reinforcement learning, and the continued advancement of the field. We are also known for our strong shared values of deploying AI in a way that brings positive benefits to society. In these turbulent times, with so much geopolitical, economic and climatic instability, the world needs more Canadian AI" (CIFAR, 2020a, p. 4).

Conclusion: Artificial Ecosystems

The awarding of the PCAIS in 2017 and its renewal in 2022 thus served as highly visible evidence of the successful positioning of CIFAR at the centre of AI research, funding, and development in Canada. The PCAIS and other initiatives have helped to popularize the idea of a Canadian "AI ecosystem" (Innovation, Science and Economic Development Canada, 2021).

But Canadian AI research and industry is more like a botanical garden than a naturally occurring environment.

Occurring mere weeks before the PCAIS announcement, a less-heralded event speaks volumes about the institute's role within Canada's cozy public-private AI garden: CIFAR moved their headquarters from the University of Toronto to the city's MaRS Discovery District.⁷ The institute themselves described this physical relocation from a university campus to a start-up incubator in revolutionary terms:

In February 2017 we relocated our offices to the MaRS Discovery District in Toronto, a premier innovation hub and home to some of Canada's most forward-looking organizations. This move was designed to transform how we think about ourselves, how we work together and how others see us.
(CIFAR, 2017b, p. 3)

MaRS is a nonprofit charitable organization whose mission is to support and advise technology start-ups, or as they put it on their website: "From advisory services to connections to talent, capital, customers and more, MaRS offers a range of high-value services that help high-growth companies succeed. We also provide access to the MaRS ecosystem, a curated community of entrepreneurs, investors, corporates, academics and government partners" (MaRS, n.d.). MaRS somewhat resembles the famous start-up accelerators of Silicon Valley (e.g., Y Combinator), but, like CIFAR, it owes its existence to substantial governmental largesse (McIntyre, 2018)—another "ecosystem" that wouldn't exist without plenty of publicly funded fertilizer.

Such exotic hybrids flourish in the artificial environment created by the confluence of neoliberalized academic research agendas and public funding directed toward private profit. CIFAR's board chair William L. Young also casts the institute's function in ecological terms: "CIFAR is an essential part of the global research and innovation landscape, filling a unique niche that complements the work of universities, governments, and the private sector. CIFAR convenes extraordinary minds with the intellectual weight of the world's top universities and

⁷ While the symbolic significance of this move is hard to ignore, note that UT helped to fund MaRS and had their own space in the MaRS building before opening the Schwartz Reisman Institute for Technology and Society across the street.

the agility and ingenuity of a start-up” (CIFAR, 2020b, p. 2). CIFAR’s relocation to a start-up incubator thus merely reflects the institute’s current mission.

These material changes speak to a more symbolic change for CIFAR: rather than fostering networks of researchers—i.e., facilitating connections between the nodes—when it comes to AI in Canada, CIFAR now occupies the network’s centre, acting as the hub that links together academia and its researchers, governments and their policy makers, and established corporations and long-shot start-ups. This move to the hub was achieved by transforming and scaling up their founding vision, mutating the ideal of unpredictably productive “research networks” into an ideological commitment to lucrative public-private partnerships and commercialized research while redefining AI to exclusively mean neural networks and the techniques of deep learning.

Chapter 6: Chinook's Nested Black Boxes by Nicolas Chartier-Edwards, Marek Blottière, and Sophie Toupin (en français)

Mr. Omorotionmwan: Do you not think . . . do you not attribute the high rate of refusal to the introduction of the Chinook model?

...

MR. GIBSON: Insofar as that's seeking an opinion, I would object and advise the witness not to answer that question.

...

Mr. Omorotionmwan: That . . . that is not an opinion. That is based on facts, the facts that we have just looked at.

Abigail Ocran vs. the Minister of Citizenship and Immigration

Introduction

En 2016–2017, le ministère d'Immigration, réfugiés et citoyenneté Canada (IRCC) a entrepris un virage computationnel pour améliorer la productivité de ses employés (*Abigail Ocran vs. the Minister of Citizenship and Immigration*, 2021). Déjà à cette époque, les délais de traitement des demandes de visa de résidence temporaires étaient lourds et plusieurs gouvernements (Américain, Européen) avaient emboîté le pas de l'automatisation des processus de traitement d'immigration. Face à cet enjeu et cette tendance, le ministère d'Immigration, réfugiés et citoyenneté Canada (IRCC) a mis en place en 2018 un nouvel outil technologique afin d'accélérer le traitement des demandes. Il s'agit de l'outil Chinook (Champagne, 2021; Cour fédérale canadienne, 2021), l'un des multiples systèmes technologiques mis en place par l'IRCC dans les dernières années. La focale de notre étude est Chinook, seul système jusqu'à ce jour ayant suscité la controverse dans les médias.

Cette réaction de l'IRCC d'adopter le virage computationnel, c'est-à-dire un retranchement immédiat dans les solutions technologiques, s'inscrit dans un paradigme que l'on peut qualifier de techno-solutionniste (Morozov, 2013), soit l'idée selon laquelle tout problème peut être résolu par une solution technique—idée largement documentée par la littérature scientifique et ce, de manière transdisciplinaire. Toutefois, cette littérature démontre que l'une des conséquences de l'adoption de ce paradigme est l'incapacité d'imaginer des solutions qui ne sont pas strictement techniciennes lorsque les technologies ne sont pas à la hauteur de leur promesse. En résulte donc une forme de dépendance qui invisibilise un ensemble d'enjeux fondamentaux auxquels la technologie sous-tend. D'ailleurs, après l'intégration de Chinook par l'IRCC, le taux de refus déjà élevé de demandes provenant de pays d'Afrique francophone a "bondi" selon certains médias, "frôlant les 100%" de refus selon Radio-Canada (Champagne, 2021; Schué, 2021). Ce type d'échec, souvent répétés, des solutions technologiques à des problèmes complexes, permettent l'observation et l'étude de controverses qui émergent dans la société. Ce rapport s'intéresse justement au cas "Chinook", à la fois comme objet technologique et comme controverse médiatique.

Dans ce rapport, nous rendrons compte des principaux aspects de cette controverse, notamment des enjeux de l'opacité autour de la technologie Chinook, de la manière dont son usage semble avoir reproduit certains biais racistes ayant suscité différentes réactions médiatiques et finalement, sur les conséquences du déploiement de technologies qui automatisent certaines fonctions gouvernementales et ce, même à un degré minime, enclenchant un processus de mutation des pratiques de gouvernance politique.

Méthodologie

La présente méthodologie s'appuie sur l'analyse des controverses, une perspective développée par des chercheurs en science et technologie (STS) pour étudier les sociétés qui façonnent et sont façonnées par les STS (Latour, 2005). Autour de l'initiative Macospol au MediaLab dirigé par Bruno Latour, Tommasino Venturini définit la controverse comme « every bit of science and technology which is not yet stabilized, closed or "black boxed." . . . We use it as a general term to describe shared uncertainty" (quoted in Venturini, 2010, p. 260). Pour Venturini, cette

incertitude partagée fait référence aux situations de discorde, de débat ou de tension qui émergent entre une pluralité d'acteurs à propos de l'actualité de la vie sociale. Les controverses ayant atteint la clôture ou la fermeture comme Max Liboiron (2021) nous le rappellent dans *Pollution Is Colonialism* peuvent être rouvertes. Une controverse n'est jamais fermée pour de bon.

Dans ce rapport, nous accordons une attention particulière aux contextes qui ont accompagné une situation de controverse. Cette attention est inspirée de la discussion de Donna Haraway (1988) sur la connaissance située ou "situated knowledge" qui invite à être attentif aux "politics and epistemologies of locating, positioning, and situating [as well as] the connections and unexpected openings [that] situated knowledges make possible" (pp. 589–90; our emphasis). Pour Haraway, le savoir est un processus politique imbriqué dans des relations de pouvoir qu'il est important de reconnaître et situer. Transposé à notre cas, la connaissance située est une manière de problématiser les relations de pouvoir rendues visibles lors de la controverse.

De la même manière, la connaissance située est une invitation à la réflexivité quant aux savoirs présentés dans ce rapport. Aussi, il est important de noter que cette recherche a été menée par une équipe de chercheur.euses évoluant dans différentes universités francophones et anglophones d'Amérique du Nord, influençant les réflexions et cadres théoriques utilisés.

En suivant cette discussion, nous avons été attentifs, à chaque étape de la recherche, aux moments d'incertitudes partagées ayant mené à des ouvertures inattendues dans le déroulement de la controverse observée. Ces différents moments ont été documentés à travers deux méthodes de recherche qualitative à savoir a) une recherche documentaire et b) des entrevues semi-dirigées.

- a) Une première recherche documentaire sur Chinook a été réalisée à l'hiver 2022 par Nicolas Chartier-Edwards. La méthode « boule de neige » a été utilisée pour la recherche de document via les outils de recherches Google, Google Scholar, Sophia (moteur de recherche de la bibliothèque de l'Université Laval) ainsi que Sociological Abstract. Certains documents ont également été obtenus par le biais de l'avocat en immigration, M. Wei William Tao, avec lequel l'une des membres de nombre équipe Sophie

Toupin est en contact via Twitter et courriel, et qui fera l'objet d'une discussion dans ce rapport. Les différents documents recueillis donnent une vue d'ensemble sur les modes de gouvernance employés (lois et règlements, déclarations, directives, lignes directrices éthiques, documents stratégiques, initiatives de recherche et financements publics majeurs), les secteurs, thèmes et problèmes perçus, et les acteurs qui ont gagné ou perdu en visibilité à travers ces processus politiques.

- b) Les entrevues semi-dirigées ont été réalisées à l'été et l'automne 2022 par Marek Blotière, Nicolas Chartier-Edward et Sophie Toupin. Notre échantillon a été construit collectivement en suivant deux considérations : (1) notre liste devait suivre un équilibre de parité et de diversité et (2) les candidat.e.s devaient avoir participé à la controverse entourant le déploiement de Chinook par l'IRCC. L'échantillon a été divisé en différents groupes de professions (avocat d'immigration, ONG, média, corporatif, académique et gouvernement). Nous avons contacté une dizaine de personnes mais la plupart ont décliné. Les entrevues ont été réalisées en ligne et enregistrées via la plateforme Zoom pour une durée de 60 minutes. Une grille de question a été réalisée autour de trois thèmes (biographique, controverse et silence).

Objet de recherche : Chinook

Cette section du rapport s'intéresse à la manière dont l'IRCC utilise de nouvelles technologies numériques ou informationnelles, qu'elles soient algorithmiques, d'analytique avancées ou d'intelligence artificielle. Ces technologies ont comme dénominateur commun le fait qu'elles sont processuelles, au sens où les possibilités d'optimisation sont constantes ainsi que le paradigme du design les font exister non comme objet ou produit fini mais comme « démo » ou « beta » de la prochaine version à venir (Amoore, 2023; Cour fédérale canadienne, 2021; Halpern & Günel, 2017). D'ailleurs, le constant prototypage de Chinook est bien illustré par les mots de Andie M. Daponte, directeur de l'optimisation et de la modernisation du réseau international à l'IRCC :

In May 2021, limited rollout of a version subsequent to the Chinook July 2020 version commenced at certain migration offices. The Chinook May 2021 version has moved from reliance on Microsoft Excel to reliance on an online platform. IRCC intends to retain all functionality from the Chinook July 2020 version in the Chinook May 2021 version. (Cour fédérale canadienne, 2021, p. 13)

L'usage de ces nouvelles technologies dans le cadre de l'exercice du politique résulte nécessairement en un réagencement des pratiques de gouvernance. Ces dernières, autrefois fondées sur l'application d'ensembles de règles, transitionnent progressivement par les algorithmes de deep learning vers la production de nouvelles règles (Amoore, 2023; Amoore & Raley, 2017). Cette ontologie processuelle constitue un puissant argument de dédouanage lorsque ces technologies ne sont pas à la hauteur de leur promesse et produisent des effets délétères réels. La faute ne résiderait jamais dans l'usage des technologies en tant que telles, mais plutôt dans le fait qu'elles n'ont pas encore été assez optimisées.

L'objet d'étude concret de ce rapport est l'outil informatique Chinook. Il s'agit d'un outil informatique qui conjugue une interface de prise de décision avec certaines technologies algorithmiques qui permettent de traiter des ensembles de dossiers. Chinook interagit avec au moins un autre outil utilisé par l'IRCC, le Global Case Management System. Il s'agit de la base de données dans laquelle les applications des demandeurs sont stockées à l'échelle globale. L'outil Chinook est utilisé depuis mars 2018 et a été « optimisé » plusieurs fois depuis (Cour fédérale canadienne, 2021). Le rapport public 8 intitulé Traitement différentiel dans le recrutement et les taux d'acceptation des étudiants étrangers au Québec et dans le reste du Canada, produit par le Comité permanent de la citoyenneté et de l'immigration (CIMM) suite à un ensemble de consultations publiques, nous fournit la définition la plus claire, publiquement disponible, de ce qu'est Chinook. Tout d'abord, Chinook n'est pas une technologie de prise de décision. Ces dernières, quant à l'acceptation ou au refus des demandes, sont prises par les agent-es d'IRCC. L'IRCC affirme que, selon ses définitions, Chinook n'est pas un programme d'analytique avancée, car « il ne fait pas de prédictions basées sur des données » (Comité permanent de la citoyenneté et de l'immigration, 2022, p.

71). Les différents mots-clés et indicateurs utilisés par le logiciel sont fabriqués et approuvés par des humains, par l'entremise de formulaires. Il n'y aurait pas non plus de formes d'apprentissage automatique à l'œuvre dans Chinook, toujours selon IRCC. Le Rapport 8 conclut donc que

On peut dire que Chinook n'est pas une forme d'intelligence artificielle, même dans le sens le plus simple de l'expression. Bien qu'il effectue une fonction de tri qui serait habituellement faite par des humains, cette fonction ne nécessite pas un grand effort intellectuel. Ses fonctions d'affichage se rapprochent davantage de celles d'un outil électronique, selon la définition d'IRCC, puisque Chinook exécute automatiquement une tâche répétitive, comme le ferait un programme de remplissage automatique de documents PDF. (p. 72)

Toutefois, le comité CIMM reconnaît que l'arrangement des documents, indicateurs et mots-clés, disposés dans une interface similaire à une feuille de calcul, cadrent les demandes, ce qui risque de « mettre en évidence certaines parties des demandes et en laisser d'autres de côté » (p. 72). La reconnaissance de ce cadrage est importante, puisque même si Chinook n'est pas une IA ou un système de prise de décision assistée, les technologies d'interface conditionnent certainement l'action de l'utilisateur (Bratton, 2015).

Par contre, le langage utilisé par l'IRCC demeure flou tout comme leurs nombreux outils technologiques qui sont difficiles à différencier. Dans un communiqué datant du 24 janvier 2022, l'IRCC mentionne l'utilisation de l'analyse avancée des données pour aider les agents d'IRCC à trier les demandes de visa de résidents temporaires et en faciliter leur traitement (Immigration, 2022). Ces systèmes automatisés comme on peut le lire ne sont pas utilisés pour rejeter des demandes comme cela est fait dans d'autres pays (Koulish & Evans, 2021), mais une technologie d'automatisation et d'aide à la prise de décision est à l'œuvre. IRCC dit clairement qu'il « s'efforce de créer et de mettre en œuvre des technologies axées sur les données de façon responsable » (Immigration, 2022). Cette manière de faire est d'ailleurs en adéquation avec les principes d'IA responsable (*La Déclaration de Montréal IA responsable*, 2018). Chinook n'est peut-être pas une IA comprise comme une technologie d'apprentissage machine, mais l'IRCC est clairement en train d'utiliser la prise de décision assistée pour améliorer ses processus que

ce soit avec ou sans Chinook. Les termes et la non différenciation des différents outils utilisés et créés par l'IRCC continuent à mettre une ombre au tableau.

Controverse : Mutation de la gouvernance

L'usage des technologies numériques et informationnelles bouleversent radicalement les modalités d'exercice de la gouvernance. Alors que cette dernière s'était bâtie au 18^e siècle sur l'étude des populations afin de produire des instruments statistiques dont la fonction était la production de la norme par l'application de règles (Foucault, 2004), nous assistons désormais à l'émergence d'une gouvernance dont l'administration des individus s'effectue en fonction de la dichotomie accès/rejet (Deleuze, 1990). Plus précisément, il s'agit d'une gouvernance dont la finalité réside dans la modulation constante, à la fois des conduites individuelles par l'action sur leur environnement d'opération mais aussi, par la modulation des instruments et technologies de gouvernances elle-même (Deleuze, 1990). Plus précisément, nous passons d'un modèle de la gouvernance qui visait la gestion du risque à l'anticipation de l'incertitude (Halpern & Günel, 2017).

En immigration, l'émergence de ce changement de paradigme, qui va de paire avec le paradigme techno-solutionniste, se produit dans le contexte de la sécurité nationale, plus précisément, de la « guerre contre la terreur » (Amoore & Raley, 2017). Les outils algorithmiques et les technologies d'intelligence artificielle deviennent des instruments de sécurisation privilégiés des états. Toutefois, nous nous trouvons ici en face d'une contradiction majeure. Simondon soulignait déjà que les technologies, dont la finalité était l'automatisation d'une action, constituaient un degré de perfectionnement technique assez pauvre. Elles nécessitent toujours plus d'informations et sont incapables d'achever le stade de « l'individuation technique », c'est-à-dire la capacité à opérer avec la plus grande ouverture dans un ensemble de milieux différents (Simondon, 2012). Puisque les technologies algorithmiques nécessitent constamment d'être abreuvées en data afin de parfaire leur fonctionnement (Cardon, 2015), elles imposent aux usager.e.s de produire des situations d'incertitudes (Amoore, 2023). L'usage de ces technologies par les États se résout donc dans l'acceptation cynique que des erreurs seront commises car les algorithmes s'appuient justement sur ces

erreurs afin de se « parfaire ». Dans le cas de Chinook, suite à quatre années d'opération au cours desquelles l'outil, en dépit de sa nature, s'est retrouvé pris dans une controverse, ce dernier sera destiné à être remplacé, malgré les mises à jour, par une éventuelle plateforme (Immigration, Refugees and Citizenship Canada, 2021). Cette plateforme est actuellement développée par les grandes firmes de consultation telles que Accenture, McKinsey et cie, Deloitte et Gartner. Celles-ci non pas pour objectif de développer l'IA mais plutôt de développer le contenant dans lequel l'analytique avancée pourra opérer. L'avocat en immigration Mario Bellissimo nous partage son analyse:

We are in a data mining process right now. This increased digitization is creating the framework for applications like Chinook to be completely collapsed into AI in the next few years. These platforms will be phased out. They are feather systems for collecting data, creating paradigms under which officers should render decisions which will likely in my view feed future algorithms. (Mario Bellissimo, personal communication, October 13, 2022)

Il renchérit en indiquant ceci:

We don't know the paradigm in which these systems are developed. (Mario Bellissimo, personal communication, October 13, 2022)

Bien que nous ne puissions pas savoir avec exactitude le type de paradigme dans lequel s'inscrit le ou les systèmes développés par l'IRCC, Petra Molnar (communication personnelle, Novembre 9, 2022) articule ce que fait l'IRCC à travers le paradigme du techno-solutionnisme. La gouvernance qui vise la gestion des incertitudes qui souscrit par le fait même au paradigme techno-solutionniste constitue une forme de cybergouvernance. «Le terme « cybergouvernance » nous semble approprié pour décrire ce nouveau paradigme puisqu'il recoupe l'usage des technologies informationnelles et numériques ainsi que les pratiques de gouvernances reposant sur les boucles de rétroaction cybernétique, entendues comme réinsertion du général sur le particulier (Freitag, 2013). Dans le cas de Chinook, il s'agit d'administrer l'accès ou le rejet à l'immigration d'un individu particulier à travers un ensemble de dossier dont la « lecture » est

effectuée par traitement algorithmique en fonction d'indicateurs généraux de « risques » en fonction des pays d'origines. Ces indicateurs s'appuient tout d'abord sur une conception abstraite des pays en question et, puisqu'ils sont produits par des humains, sont susceptibles de contenir des biais racistes. D'ailleurs, l'IRCC a dévoilé en 2022 une stratégie anti-raciste 2.0 accompagnée d'une approche d'égalité par design (equality-by-design) comme résultante de la controverse entourant les taux de rejet des étudiant.e.s africain.e.s.

Controverse: Chinook en tant que “Black-Box”

Le sujet de la boîte noire (ou black-box en anglais), est, depuis 2016 et la parution de *The Black Box Society* de Frank Pasquale, un enjeu qui a gagné en visibilité. En effet, le fonctionnement opaque des plateformes digitales comme Facebook ou Google ont contribué à mettre sur le devant de l'agenda politique les questions de transparence (transparency) et de responsabilité (accountability) (Ananny & Crawford, 2016). L'usage répété du concept de boîte noire, notamment dans les sphères médiatiques, a cependant contribué à en réduire le sens, donnant parfois le sentiment de “perte de contrôle” face à l'automatisation de certaines tâches (Bucher, 2016). Inversement, la rhétorique consistant à promouvoir la transparence s'inscrit dans une volonté de retrouver un pouvoir d'agir sur les technologies (voir Roberge et al., 2019). Cette tension entre opacité et transparence est au cœur des enjeux soulevés par la controverse de Chinook.

Dans la sphère juridique, les interrogations autour de Chinook comme boîte-noire s'articulent autour de deux éléments à savoir a) le lien entre Chinook, l'IA et les autres outils d'analyse avancée b) le manque de transparence d'IRCC face au déploiement des nouvelles technologies et comment elles interagissent ensemble. La controverse autour du lien entre Chinook et IA met en scène l'avocat Wei William Tao qui argumente que Chinook est une IA et l'IRCC qui dément cette affirmation. L'opacité de Chinook est critiquée par Mr. Tao qui dénonce certaines contradictions dans le fonctionnement de l'outil Chinook comme le fait que ce dernier assiste algorithmiquement les fonctionnaires dans l'exercice de leurs fonctions. Cette critique rejoint directement le second élément puisqu'elle démontre le manque de transparence de l'IRCC (à

ce moment-là de la controverse) dans l'explication du fonctionnement de Chinook. Sur ce point, l'IRCC suggère à ses fonctionnaires, via un manuel d'instruction sur Chinook, de "limiter l'information partagée sur Chinook afin d'empêcher les applicant.e.s de découvrir des angles morts dans la programmation, leur permettant de tricher" (Chartier-Edwards, 2022). Comme nous allons voir à présent, ce dernier élément est problématique pour les journalistes qui ont couvert ce sujet.

Dans la sphère médiatique, "C'est une affaire opposant ce ministère fédéral [IRCC] et Abigail Ocran, une étudiante du Ghana refusée, qui a forcé à ouvrir une première fenêtre sur Chinook, dont on ignorait l'existence jusqu'à maintenant" (Champagne, 2021). Le déploiement de Chinook par l'IRCC coïncide d'ailleurs avec l'accroissement du nombre de rejets d'étudiant.e.s en provenance de pays d'Afrique francophone selon la presse québécoise. Pour la journaliste Sarah Champagne, le taux de refus pose un certain nombre de problèmes notamment de biais racistes dont l'origine est difficile à retracer. Lorsqu'elle couvre le sujet, la journaliste en immigration se heurte à des problèmes d'accès à l'information (Sarah Champagne, communication personnelle, 8 août 2022). Ce problème est partagé par un autre interlocuteur qui dirige une association de défense des droits des immigrants et qui compare l'IRCC au Service Canadien du Renseignement de Sécurité (SCRS) dans lequel les fonctionnaires ont des obligations de confidentialité les "interdisant de discuter dans les couloirs de ce type de dossier" (Thibault Camara, communication personnelle, 19 septembre 2022). Notre équipe de recherche s'est aussi butée au long délais pour obtenir des demandes d'accès à l'information qui nous reviennent caviardées. L'opacité entourant Chinook est un enjeu politique qui formalise certaines pratiques et processus tenus secrets pour des raisons de sécurité nationale. D'ailleurs, la plupart des interlocuteurs que la journaliste a contactée ont principalement témoigné "off the record" ou "on background", contraignant la journaliste à ne pas divulguer certaines informations ou à rester vague quant à l'explication de l'outil Chinook.

De plus, la journaliste Sarah Champagne constate qu'aucune considération juridique n'avait été prise en compte lors du déploiement initial de Chinook et n'avait pas fait l'objet d'évaluation ou de commission indépendante. Ce constat est partagé par un autre interlocuteur, Thibault Camara, qui dénonce l'absence de charte de développement et un

retard dans la régulation de l'innovation dans le domaine public. Petra Molnar, de son côté, dénonce le fait que le Canada n'a pas de charte des droits humains à respecter dans le domaine de l'innovation technologique (Petra Molnar, communication personnelle, 8 novembre 2022). Aussi bien du côté de la critique juridique que de la critique médiatique, l'opacité est dénoncée principalement sous l'angle d'un certain racisme systémique présent à l'IRCC qui favoriserait des discours et pratiques en vase-clos. Le manque de transparence de Chinook a été justifié par l'IRCC de telle manière qu'un flou dans son fonctionnement et ses finalités a été maintenu. L'opacité comme stratégie devient dès lors un argument politique de protection non seulement des institutions canadiennes mais plus largement d'une identité nationale canadienne.

Controverse: Racisme systémique au sein de l'IRCC

Selon Camara, un acteur de la défense des droits des immigrants, la question du racisme systémique dans les institutions nord-américaines, et plus particulièrement dans les ministères d'immigration, est un enjeu récurrent dans la critique des processus déshumanisants que sont les demandes d'immigration. Chinook est en ce sens l'incarnation, ou la formalisation, de ce processus déjà existant qui catégorise et classifie des demandeurs de visas selon un certain nombre de critères et de variables, réduisant et simplifiant la réalité complexe d'individus issus de différents contextes socio-culturels. Pour le milieu associatif, Chinook, en tant qu'objet de nouveauté associée notamment à des innovations comme l'IA, a été une opportunité pour rendre visible ces enjeux. Comme le mentionne cet intervenant,

«Si, pour la énième fois je répète, « Ah le processus d'immigration est ralenti parce que Québec et Ottawa ne sont pas d'accord sur les seuils », au bout de 15 fois, le journaliste ne va plus l'écrire ce papier. [Alors on est allé] chercher un « buzz word » que personne ne connaissait. Donc on est allé demander l'avis à des experts et là on a ouvert tout un nouveau volet [à travers la notion d'Intelligence Artificielle]. Puis [...] nous on savait, moi je savais ce que c'est l'IA [et ce qui ne l'était pas], mais juste le *mot clé* [intelligence artificielle], même si c'était pas vraiment ça, ça nous a permis de faire du millage là-dessus ». (Thibault Camara, communication personnelle, 19 septembre 2022)

Dans la sphère médiatique canadienne et québécoise, les nouvelles technologies comme l'IA ont fait l'objet d'un engouement médiatique permettant de mobiliser des audiences et rendre visible des enjeux (Blottière, 2023).

Bien que cette visibilité soit problématique comme nous allons le voir dans le prochain paragraphe, elle est nécessaire dans un contexte où les demandeurs de visas qui ont vu la demande être refusée n'étaient ni citoyens, ni physiquement sur le territoire canadien et souvent n'avait pas ou peu de contacte sur place. Ce faisant, la possibilité de rendre visible leurs situations était, dans le contexte opaque que représente les processus d'immigration, quasi nulle. À ce sujet, différentes stratégies de mobilisation ont été mises en place par les demandeurs de visa comme l'envoi massif de lettres à des journalistes provenant d'étudiant.e.s indien.n.e.s (*Sold a Lie*, 2022) ou encore la mobilisation des familles de demandeurs résidant au Canada. Du côté des militants de défense des droits, l'enjeu de visibilité est aussi important afin de mettre à l'agenda politique les questions de racisme et d'atteinte aux droits des migrants (Voirol, 2005).

En s'emparant de cette controverse, il semblerait que le cycle des nouvelles médiatiques a transféré l'enjeu du racisme systémique des demandes issus de pays d'Afrique à une question de politique identitaire anti-francophone au Canada. Plus particulièrement, la question de la régulation des processus d'automatisation des services publics que soulève pourtant Chinook, ou encore, de celle sous-jacente à cette dynamique de normalisation et de formalisation de certains biais raciste ont été, au fur et à mesure que la controverse a avancé, remplacé par un discours sur la discrimination linguistique, notamment au Québec. Ce glissement médiatique est problématique du fait de l'invisibilisation, ou le déni, du racisme systémique présent institutionnellement au Québec et de la réappropriation de ce sujet à des fins de politiques identitaires en témoigne un article du *Nouvelliste* intitulé "le Bloc [québécois] le taux de refus des étudiants francophones étrangers". L'emphase est mise sur l'aspect francophone déplaçant ainsi la question du racisme systémique présent dans le système interculturel québécois (voir Salée, 2010) à une dénonciation d'une discrimination linguistique orchestrée par le Canada.

Conclusion

Cette brève étude de cas fait état d'un virage computationnel au sein de l'État canadien. L'automatisation des processus bureaucratiques fait grâce à la datafication, la plateformes et la frontière numérique est en fait une tendance EuroAméricaine. La croyance que l'innovation technologique va non seulement améliorer l'efficacité des bureaucraties d'une part et d'autre part résoudre les problèmes d'inégalité sociale, économique et environnementale est au cœur du problème. Les absents de cette recherche sont celles et ceux dont les visas ont été rejetés pour différentes raisons.

Appendix by Sophie Toupin and Fenwick McKelvey

How do you study a policy process that was not a process? AI governance in Canada is immediately incoherent. Cities, provinces, and the many branches of the federal government have participated separately and collaboratively in drafting what we might call AI governance. The result, to borrow a phrase from policy scholar Tess Lea, is “vertiginous and incoherent” (Lea, 2020a, p. 25). That incoherence can be lost in much of the AI governance literature that either focuses on national AI approaches without the depth to consider inconsistent internal updates (Radu, 2021), or on definitions of good AI governance (Floridi & Cowls, 2019; Margetts, 2022; Mittelstadt et al., 2016; Ulnicane et al., 2021). What we have instead is a mess and our methods reflect this challenge.

Controversy studies is the usual way to map a messy policy process, but that too proved difficult. If “controversies are situations where actors disagree (or better, agree on their disagreement)” (Venturini, 2010, p. 261), then Canada might be seen as a cold controversy, one marked by “reciprocal indifference” (Venturini, 2010, p. 264). But we found more of a pattern of one-sided institutional indifference marked by a willingness to hear disagreements that ultimately do not matter much. Canada habitually holds ad-hoc consultations and symbolic parliamentary hearings (Webber et al., 2023). These hearings can be open or closed, but they are much narrower than possible consultations that allow for contributions that contest the terms of reference and are generally indifferent to the responses. (see Binderkrantz et al., 2021). More accurately, even when summarized in a “What We Heard Report: 2022 Roundtables on Online Safety,” the actual policy recommendations ignore anything disagreeable that they heard.

The challenge became the cool situation of AI governance in Canada. A cool situation refers to issues that have become accepted as uncontroversial—a result we explained as a product of scientist-entrepreneurs, governments, and journalists working to *freeze out* AI’s controversiality in the name of supporting Canada’s AI ecosystem. Dandurand, McKelvey, and Roberge write, “Freezing out cools down potentially hot situations through interactions between key actors and institutions that create a mutually beneficial framing of the said situation—resulting in what looks like a cool situation” (Dandurand et al., 2023, p. 3) As we

discuss in Chapter 4, public consultation around AI largely functions in the same way to close debate and to maintain the state's legitimacy to regulate AI. How might we study a lack of controversy? The implications require balancing between methods based on an empiricism of presence with an empiricism of absence. Following Marres (2020), we match controversy mapping with situational analysis, using the latter to fill in the former. The result turns from controversy analysis as interpretative to constructivist and normative, allowing us to question what should be controversial but is not.

Our approach moved from controversy study to defining AI's situation, building on the feminist and anti-oppressive methodology called situational analysis. A situation is "both an object confronted and an ongoing process subsequent to that confrontation" (A. E. Clarke, 2005, p. 21). As a methodology, this approach seeks to analyze a particular situation of interest through the specification, representation, and subsequent examination of the most salient elements in that situation and their relations.

To put situational analysis in practice, we first asked: What is policy? What is a document? We considered a policy to be a paper lodestone. Policy sets objectives and goals that, in the strictest sense, bind actors to interpret and actualize them. The interpretation of policy is then the work of the policy process, translating policy into regulation (Hunt & McKelvey, 2019). We can identify at least four modes of policy agency:

1. Government policy in laws or bills that have the backing of the state to actualize its societal goals and objectives;
2. Corporate or organizational policies embedded into governing documents and attached to jobs and accountability mechanisms (e.g., community standards);
3. Governmental or organizational policy that defines accountabilities or functions for key roles or personnel (e.g., mandate letters); and,
4. All sorts of speculative or aspirational policy that exists without a clear instrument of realization.

In Canada, there is an absence of AI-specific government policy in law. There may, however, be ways that AI becomes a policy instrument without being stated as such. The Canadian

government's proposed Online Harms bill included a twenty-four-hour takedown requirement that acknowledged that AI could be a regulatory instrument, but did not define it.

We attended to policy domains in our project by:

1. Focusing on policy documents aiming to shape the development of AI (in keeping with our inter-team obligations) while acknowledging that Canadian jurisdictions have not established any new laws for AI; and,
2. Selecting cases of known controversies where AI was invoked as a policy instrument, but where we did not know beforehand that AI was considered.

We wish to acknowledge that "AI policy" is a limited account of AI's governmentality and that other facets of our research team will consider the applied nature of AI without the policy referentiality mandated by the ORA project.

Policy has instruments to actualize and interpret its commands. We should note that an AI interprets policy differently from a human, but these are wider considerations. For our project, we coded policy documents based on whether:

1. they were attached to funding commitments;
2. they established new organizations;
3. they designated or changed the mandate of existing organizations; or
4. there were other policy documents that stabilized the one being coded.

We left our definitions loose, looking at a range of documents that make a claim to be a policy document. Documents "are not simply instruments of bureaucratic organizations, but rather are constitutive of bureaucratic rules, ideologies, knowledge, practices, subjectivities, objects, outcomes, and even the organizations themselves" (Hull, 2012, p. 253). And yet, documents arrive extracted from means of studying their construction at a pace that exceeds close reading. A comparable study found 84 AI governance documents in Canada produced by governments alone (Attard-Frost et al., 2023). If we stay true to the porous nature of AI governance, then this number only grows when including civil society and industry.

Then, we collected governmental policy documents (federal and provincial) where AI is mentioned between 2012 and 2021. The documents we collected represent a broad and purposeful overview of policy documents found on governmental (federal, Quebec, and Ontario) and nonprofit organizations' websites. We found 300 documents that corresponded to this description. Using the reference management program Zotero as a tool for situational analysis, we identified trends in who produces AI policy documents. Through the collection process, we developed a sense of the key or most influential documents that required further investigation. We selected 30 documents that best reflected the key actors and influences in our situation. Importantly, we selected documents that went through a process of consultation.

As we can see from the map below, most of the policy documents on AI came out of Ottawa, i.e., from the federal government, and from the three provinces considered to be AI hubs (Alberta, Ontario, and Quebec).



Image 2: A representation of the location of the policy documents

Documents analyzed by type




Image 3: The majority of the policy documents were found on governmental webpages

Government Documents about AI

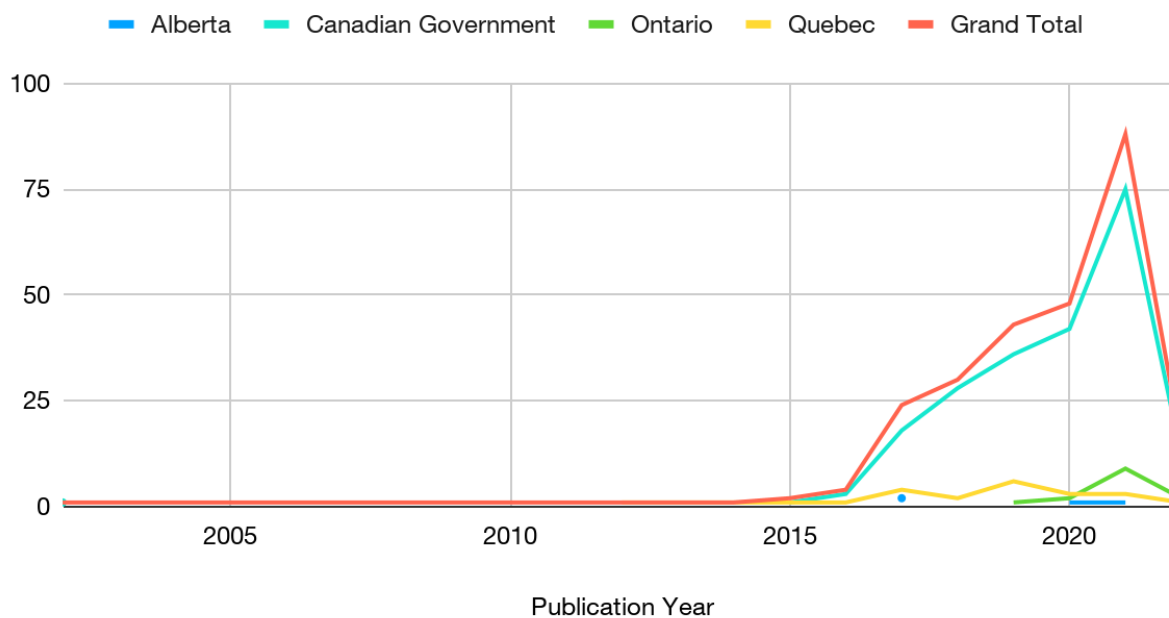


Image 4: Grouped by publication year, this graph shows that the documents started appearing in 2016–2017 with a peak in 2020–2021.

Inspired by materialist and research-creation approaches to the study of assemblages, we organized three participatory workshops with our research team at Concordia University and INRS. The first workshop aimed to create a messy map of the main topics addressed in 30 documents. The second workshop was a transition from a messy map to a coding strategy. The final workshop allowed us to pick five case studies to understand AI governance controversies in Canada.

Workshop 1

We organized a workshop to make sense of the 30 policy documents we had collected through desktop research. Prior to the workshop held on 21 March 2022, participants were assigned to read five to seven policy documents each out of the 30 selected. The goal was to quickly map the ecosystem around the policy documents, make them relate to one another, and connect them to local, national, and global trends. We were interested in collectively thinking through

the controversies and closures surrounding our list of policy documents and identifying driving actors, forces, and dynamics. The participants were asked to carefully read the assigned policy documents and consider the following questions:

1. What is the date of the document? What kind of document is it? (e.g., white paper, report, statement, brief, memo, bill)?
2. Are there funding commitments attached to the policy document? Does it establish a new organization? If so, which one? Or to which organization(s) is it tied? Does it designate or change the mandate of existing organizations? If so, how? And, are there other documents that stabilize this policy document?
3. Whose work or life is affected/targeted by this document (e.g., governmental employees, academia, racialized populations)?
4. What are the actions taken in the document (e.g., establishing law, improving internal processes, funding researchers)?
5. What are the **issues addressed** in the document (e.g., economic recovery after Covid, economic prosperity, predicting crime, respecting privacy laws, ethical AI)?
6. What are the **expressed narratives** associated with the policy documents and what are its associated promises or desired futures (e.g., use of AI for efficiency, sovereignty, economic prosperity, speed, solving a crisis)?
7. What are some of the **controversies, dynamics, and forces** associated with this policy document?
8. How do they relate to the local, national, and international context?
9. How are they related to civil society or to corporations? What else is missing to understand how they are shaped and their materiality?

Out of this workshop, we produced three messy maps for three parts of the AI governance situation in Canada: Federal/Canada, Quebec/Montréal, Ontario/Toronto (see Images below). These maps deliberately mixed terms and concepts as a grounded strategy for subsequent document coding. The images below include key actors, policy controversies,

narratives, and other themes inductively drawn from the documents. We used these maps to develop a coding strategy for the second workshop.



Quebec AI policy ecosystem

Workshop 2

The transition from a messy map to a coding strategy occurred in a second workshop. We proceeded to code documents according to concepts that emerged at the first workshop:

1. Identify key actors mentioned in the document
2. Key narratives
3. Issues or concerns and examples motivating the document
4. Potential controversies or debates in the document
5. If there was an evaluation grid associated with the policy document (Yes or No)
6. If there was a consultation associated with the policy document
7. What were the silences or what might be missing in the document that is part of the larger situation?

Finally, each team member produced a short memo about their documents. We used these codes to analyze key trends in the documents presented in the next section.

Workshop 3

While we were able to provide a broad scan of trends in AI, the situation mapping prompted a motivation to focus more on specific cases that exemplified the controversies found in Canadian AI governance. Our scan had demonstrated that debates over government and police technologies had had a demonstrable impact on AI, especially the international Clearview AI scandal as well as a local scandal involving discrimination in the immigration system attributed to automated processing. In tandem, we found that the 2017 Pan-Canadian AI Strategy remained the defining policy document, one that importantly concentrated federal power and influence in the Canadian Institute for Advanced Research (CIFAR). As for actual AI policy, the most we could find were the efforts of the federal government to regulate its own use of AI. These discussions resulted in a consensus around five case studies where AI became invoked as a policy instrument in Canada. All cases expressly deal with the “reconstruction of

values and meanings in AIT-policy- interactions” that have positioned AI as an innovation needed by governments (R. Paul, 2022, p. 505).

References

- Abigail Ocran vs. the Minister of Citizenship and Immigration* (2021).
- AI Now Institute. (2018). *Algorithmic Impact Assessments: Toward Accountable Automation in Public Agencies*. <https://ainowinstitute.org/publication/algorithmic-impact-assessments-toward-accountable-automation-in-public-agencies>
- Amoore, L. (2022). Machine learning political orders. *Review of International Studies*, 1–17. <https://doi.org/10.1017/S0260210522000031>
- Amoore, L. (2023). Machine learning political orders. *Review of International Studies*, 49(1), 20–36. Cambridge Core. <https://doi.org/10.1017/S0260210522000031>
- Amoore, L., & Raley, R. (2017). Securing with algorithms: Knowledge, decision, sovereignty. *Security Dialogue*, 48(1), 3–10. <https://doi.org/10.1177/0967010616680753>
- Ananny, M., & Crawford, K. (2016). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3). <https://doi.org/10.1177/1461444816676645>
- Attard-Frost, B., Brandusescu, A., & Lyons, K. (2023). *The Governance of Artificial Intelligence in Canada: Findings and Opportunities from a Review of 84 AI Governance Initiatives* (SSRN Scholarly Paper 4414212). <https://doi.org/10.2139/ssrn.4414212>
- Axelrod, P., Desai-trilokekar, R., Shanahan, T., & Wellen, R. (2011). People, processes, and policy-making in Canadian post-secondary education, 1990-2000. *Higher Education Policy*, 24(2), 143–166. <https://doi.org/10.1057/hep.2010.29>
- Balingall, A. (2018, February 15). Ottawa awards \$950 million to five innovation ‘superclusters.’ *Thestar.Com*.

<https://www.thestar.com/news/canada/2018/02/15/ottawa-awards-950-million-to-five-innovation-superclusters.html>

Ball, D. P. (2017, March 22). Edmonton to share feds' artificial intelligence \$125M budget boost. *Metro News*.

<https://web.archive.org/web/20180313113954/http://www.metronews.ca/news/edmonton/2017/03/22/edmonton-nets-artificial-intelligence-boost-in-2017-budget.html>

Balsillie, J. (2020, June 15). *Remarks at the Industry, Science and Technology Committee* [Open Parliament]. <https://openparliament.ca/committees/industry/43-1/24/jim-balsillie-1/only/>

Bareis, J., & Katzenbach, C. (2021a). Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values*, 01622439211030007. <https://doi.org/10.1177/01622439211030007>

Bareis, J., & Katzenbach, C. (2021b). Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values*, 0162243921103000. <https://doi.org/10.1177/01622439211030007>

Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional Expressions Reconsidered: Challenges to Inferring Emotion from Human Facial Movements. *Psychological Science in the Public Interest*, 20(1), 1–68. <https://doi.org/10.1177/1529100619832930>

Beck, U. (2008). World at Risk: The New Task of Critical Theory. *Development and Society*, 37(1), 1–21.

Beer, D. (2017). The Social Power of Algorithms. *Information, Communication & Society*, 20(1), 1–13. <https://doi.org/10.1080/1369118X.2016.1216147>

- Bell, G., Blythe, M., & Sengers, P. (2005). Making by making strange: Defamiliarization and the design of domestic technologies. *ACM Transactions on Computer-Human Interaction*, 12(2), 149–173. <https://doi.org/10.1145/1067860.1067862>
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity.
- Berjis, P. (2012). Effectiveness of Prequalification Practices in Public Procurement. *Masters of Applied Sciences Department of Civil Engineering University of Toronto*, 81.
- Binderkrantz, A. S., Blom-Hansen, J., & Senninger, R. (2021). Countering bias? The EU Commission's consultation with interest groups. *Journal of European Public Policy*, 28(4), 469–488. <https://doi.org/10.1080/13501763.2020.1748095>
- Blottière, M. (2023). *Entre économie politique des controverses et visibilités controversées, la construction médiatique de l'intelligence artificielle au Canada et au Québec* [Master's, Institut national de la recherche scientifique]. <https://espace.inrs.ca/id/eprint/14207/>
- Brandusescu, A. (2021). *Artificial intelligence policy and funding in Canada: Public investments, private interests* (p. 62). Centre for Interdisciplinary Research on Montreal, McGill University.
- Bratton, B. H. (2015). *The stack: On software and sovereignty*. MIT Press.
- Brown, R. C. (2007). *A generation of excellence: A history of the Canadian Institute for Advanced Research*. University of Toronto Press.
- Bucher, T. (2016). Neither Black Nor Box: Ways of Knowing Algorithms. In *Innovative Methods in Media and Communication Research* (pp. 81–98). Palgrave Macmillan, Cham. https://doi.org/10.1007/978-3-319-40700-5_5

- Bucher, T. (2019). Bad Guys and Bag Ladies: On the Politics of Polemics and the Promise of Ambivalence. *Social Media + Society*, 5(3), 2056305119856705.
<https://doi.org/10.1177/2056305119856705>
- Calleja, D. (2001). The mother of invention. *Canadian Business*, 74(2), 56.
- Canadian Heritage. (2021, June 29). *Guiding principles on diversity of content online*.
<https://www.canada.ca/en/canadian-heritage/services/diversity-content-digital-age/guiding-principles.html>
- Canadian Heritage. (2022, February 3). *The Government of Canada Releases What We Heard: The Government's Proposed Approach to Address Harmful Content Online*.
<https://www.newswire.ca/news-releases/the-government-of-canada-releases-what-we-heard-the-government-s-proposed-approach-to-address-harmful-content-online-867463514.html>
- Cardon, D. (2015). *A quoi rêvent les algorithmes*. Seuil. <https://www.seuil.com/ouvrage/a-quoi-revent-les-algorithmes-dominique-cardon/9782021279962>
- Cardoso, T., & Curry, B. (2021, February 7). National Defence skirted federal rules in using artificial intelligence, privacy commissioner says. *The Globe and Mail*.
<https://www.theglobeandmail.com/canada/article-national-defence-skirted-federal-rules-in-using-artificial/>
- Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., & Floridi, L. (2018). Artificial Intelligence and the 'Good Society': The US, EU, and UK approach. *Science and Engineering Ethics*, 24(2), 505–528. <https://doi.org/10.1007/s11948-017-9901-7>
- Century, M. (2022). *Northern Sparks: Innovation, Technology Policy, and the Arts in Canada from Expo 67 to the Internet Age*. <https://doi.org/10.7551/mitpress/10818.001.0001>

- Champagne, S. R. (2021, November 26). Un outil informatique mis en cause dans la hausse des refus de permis d'études. *Le Devoir*.
<https://www.ledevoir.com/societe/650031/immigration-un-outil-informatique-mis-en-cause-dans-la-hausse-des-refus-de-permis-d-etudes>
- Chartier-Edwards, N. (2022). *Automatisation de l'immigration au Canada: Rapport de recherche sur le cas « Chinook »*.
- Chilvers, J., & Kearnes, M. (2020). Remaking Participation in Science and Democracy. *Science, Technology, & Human Values*, 45(3), 347–380.
<https://doi.org/10.1177/0162243919850885>
- CIFAR. (n.d.). *Presidents Emeriti*. CIFAR. Retrieved January 20, 2023, from <https://cifar.ca/leadership/presidents-emeriti/>
- CIFAR. (2009). *The exploration edition: Annual report 08–09* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2010). *CIFAR annual performance report: 2009–2010* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2012). *CIFAR annual performance report: 2011–2012* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2013). *Asking the right questions can change the world: Annual performance report 2012–2013* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2015). *Connecting the best minds for a better world: CIFAR annual performance report 2014–15* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2016). *Global momentum: Annual report 2015–16* [Annual Report]. Canadian Institute for Advanced Research.

- CIFAR. (2017a). *Canada funds \$125 million Pan-Canadian Artificial Intelligence Strategy*. Canadian Institute for Advanced Research. <https://www.newswire.ca/news-releases/canada-funds-125-million-pan-canadian-artificial-intelligence-strategy-616876434.html>
- CIFAR. (2017b). *Increasing impact: Annual report 2016–2017* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2017c, March 22). *Canada funds \$125 million Pan-Canadian Artificial Intelligence Strategy*. <https://www.newswire.ca/news-releases/canada-funds-125-million-pan-canadian-artificial-intelligence-strategy-616876434.html>
- CIFAR. (2019a). *AICan 2019: Annual Report of the CIFAR Pan-Canadian AI Strategy* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2019b). *CIFAR Annual Report 2018–2019* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2020a). *AICan 2020: CIFAR Pan-Canadian AI Strategy Impact Report*. Canadian Institute for Advanced Research.
- CIFAR. (2020b). *Impact report 2019–2020* [Annual Report]. Canadian Institute for Advanced Research.
- CIFAR. (2022). *Impact report 2021–22* [Annual Report]. Canadian Institute for Advanced Research.
- Clarke, A. (2019). *Opening the government of Canada: The federal bureaucracy in the digital age*. UBC Press.
- Clarke, A. E. (2005). *Situational analysis: Grounded theory after the postmodern turn*. Sage Publications. <http://catdir.loc.gov/catdir/toc/ecip052/2004022902.html>

- Clarke, A. E., Friese, C., & Washburn, R. (2018). *Situational analysis: Grounded theory after the interpretive turn* (Second edition). SAGE.
- Clarke, J., Bainton, D., Lendvai, N., & Stubbs, P. (2015). *Making policy move: Towards a politics of translation and assemblage* (1st ed.). Bristol University Press; JSTOR.
<https://doi.org/10.2307/j.ctt1t89fpb>
- Comité permanent de la citoyenneté et de l'immigration. (2022). *Traitement différentiel dans le recrutement et les taux d'acceptation des étudiants étrangers au Québec et dans le reste du Canada* (Rapport du Comité no 8 CIMM (44-1)). Chambre des communes du Canada. <https://www.noscommunes.ca/DocumentViewer/fr/44-1/CIMM/rapport-8>
- Copeland, E. (2018, February 20). 10 Principles for Public Sector Use of Algorithmic Decision Making. *Nesta*. <https://www.nesta.org.uk/blog/10-principles-for-public-sector-use-of-algorithmic-decision-making/>
- Cour fédérale canadienne. (2021). *Affidavit of Andie Daponte*.
- Dandurand, G., Blottière, M., Jorandon, G., Gertler, N., Wester, M., Chartier-Edward, N., Roberge, J., & McKelvey, F. (2022). *Training the news: Coverage of Canada's AI hype cycle (2012–2021)*. INRS - Urbanisation Culture Société.
<https://espace.inrs.ca/id/eprint/13149/>
- Dandurand, G., McKelvey, F., & Roberge, J. (2023). Freezing out: Legacy media's shaping of AI as a cold controversy. *Big Data & Society*, 10(2), 20539517231219242.
<https://doi.org/10.1177/20539517231219242>
- Deleuze, G. (1990). Post-Scriptum sur les sociétés de contrôles. *L'autre Journal*, 1(1).

- Deloitte. (2020). *Canada's AI imperative—Overcoming risks, building trust*. Deloitte.
<https://www2.deloitte.com/ca/en/pages/deloitte-analytics/articles/canadas-ai-imperative.html>
- Doloreux, D., & Frigon, A. (2022). The Innovation Superclusters Initiative in Canada: A new policy strategy? *Science and Public Policy*, 49(1), 148–158.
<https://doi.org/10.1093/scipol/scab071>
- Dubois, E., & McKelvey, F. (2019). Political Bots: Disrupting Canada's Democracy. *Canadian Journal of Communication*, 44(2), 27–33. <https://doi.org/10.22230/cjc.2019v44n2a3511>
- Ducas, I. (2020, September 22). Montréal: Pas de reconnaissance faciale par les policiers sans l'accord des élus. *La Presse*. <https://www.lapresse.ca/actualites/2020-09-22/montreal/pas-de-reconnaissance-faciale-par-les-policiers-sans-l-accord-des-elus.php>
- Elam, M., & Bertilsson, M. (2003). Consuming, Engaging and Confronting Science: The Emerging Dimensions of Scientific Citizenship. *European Journal of Social Theory*, 6(2), 233–251. <https://doi.org/10.1177/1368431003006002005>
- Ezeonu, I. (2009). *Gun Violence in Toronto: Perspectives from the Police* (SSRN Scholarly Paper 1586094). <https://doi.org/10.1111/j.1468-2311.2009.00603.x>
- Fahlman, S. E., & Hinton, G. E. (1987). Connectionist architectures for artificial intelligence. *Computer*, 20(1), 100–109. <https://doi.org/10.1109/MC.1987.1663364>
- Fasken (Director). (2020, October 28). *Part 1 - Introduction and Refresher: Fundamentals of Canadian Federal Procurement*. https://www.youtube.com/watch?v=tMg_8u0v8vo

- Fels, S. S., & Hinton, G. E. (1993). Glove-Talk: A neural network interface between a data-glove and a speech synthesizer. *IEEE Transactions on Neural Networks*, 4(1), 2–8.
<https://doi.org/10.1109/72.182690>
- Floridi, L., & Cowls, J. (2019). A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1), 3–71. [https://doi.org/10.1016/0010-0277\(88\)90031-5](https://doi.org/10.1016/0010-0277(88)90031-5)
- Foucault, M. (2004). *Sécurité, territoire, population: Cours au collège de France (1977-1978)*. Gallimard.
- Fourcade, M., & Gordon, J. (2020a). Learning Like a State: Statecraft in the Digital Age. *Journal of Law and Political Economy*, 1(1). <https://doi.org/10.5070/LP61150258>
- Fourcade, M., & Gordon, J. (2020b). Learning Like a State: Statecraft in the Digital Age. *Journal of Law and Political Economy*, 1(1). <https://doi.org/10.5070/LP61150258>
- Freedman, D. (2010). Media Policy Silences: The Hidden Face of Communications Decision Making. *The International Journal of Press/Politics*, 15(3), 344–361.
<https://doi.org/10.1177/1940161210368292>
- Freitag, M. (2013). *Culture, pouvoir, contrôle* (Vol. 3). Liber.
<https://www.gallimardmontreal.com/catalogue/livre/dialectique-et-societe-3-culture-pouvoir-controle-freitag-michel-9782895783909>
- Gertler, N. (2023). *Hacking AI Governance: Exploring the Democratic Potential of Canada's Algorithmic Impact Assessment* [Thesis]. Concordia University.

- Gillis, W. (2018, July 23). *Civil rights organization urges city hall to delay purchase of gunshot-location technology*. Toronto Star. <https://www.thestar.com/news/gta/2018/07/23/ccla-urges-toronto-city-hall-to-delay-purchase-of-gunshot-location-technology.html>
- Glor, E. D. (2001). Has Canada Adopted the New Public Management? *Public Management Review*, 3(1), 121–130. <https://doi.org/10.1080/14616670010009414>
- Goffey, A. (2008). Algorithm. In M. Fuller (Ed.), *Software Studies: A Lexicon* (pp. 15–20).
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Government of Canada. (2021a, May 25). *Views of Canadians on Artificial Intelligence: Final Report - Public opinion research* [Index Pages; Public Opinion Research Reports]. <https://ised-isde.canada.ca/site/public-opinion-research/en/views-canadians-artificial-intelligence-final-report>
- Government of Canada. (2021b, July 16). *A Consultation on a Modern Copyright Framework for Artificial Intelligence and the Internet of Things* [Comments]. Innovation, Science and Economic Development Canada. <https://ised-isde.canada.ca/site/strategic-policy-sector/en/marketplace-framework-policy/copyright-policy/consultation-modern-copyright-framework-artificial-intelligence-and-internet-things-0>
- Government of Canada. (2022). *Learning Together for Responsible Artificial Intelligence* (978-0-660-41696-0). Innovation, Science and Economic Development Canada. <https://ised-isde.canada.ca/site/advisory-council-artificial-intelligence/en/public-awareness-working-group>
- Government of Canada. (2024). *Responsible use of artificial intelligence (AI)*. <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai.html>

- Government of Ontario. (2021, May 6). *Consultation: Ontario's Trustworthy Artificial Intelligence (AI) Framework*. <http://www.ontario.ca/page/ontarios-trustworthy-artificial-intelligence-ai-framework-consultations>
- Gross, M. (2016). Risk as zombie category: Ulrich Beck's unfinished project of the 'non-knowledge' society. *Security Dialogue*, 47(5), 386–402.
<https://doi.org/10.1177/0967010616645020>
- Guice, J. (1998). Controversy and the state: Lord ARPA and intelligent computing. *Social Studies of Science*, 28(1), 103–138.
- Halpern, O., & Günel, G. (2017). Demoing unto Death: Smart Cities, Environment, and Preemptive Hope. *The Fibreculture Journal*, 29. <https://doi.org/10.15307/fcj.29.215.2017>
- Hannah, A., Tchilingirian, J., Botterill, L., & Attwell, K. (2023). The role of 'non-knowledge' in crisis policymaking: A proposal and agenda for future research. *Evidence & Policy*, 19(1), 116–130. <https://doi.org/10.1332/174426421X16552882375377>
- Hansen, S. S. (2021). Public AI imaginaries: How the debate on artificial intelligence was covered in Danish newspapers and magazines 1956–2021. *Nordicom Review*, 43(1), 56–78. <https://doi.org/10.2478/nor-2022-0004>
- Hao, K. (2019, June 6). *Training a single AI model can emit as much carbon as five cars in their lifetimes*. MIT Technology Review.
<https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/>
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575–599.
<https://doi.org/10.2307/3178066>

- Hemmadi, M. (2019, November 28). *Ottawa's AI ethics test didn't consider Palantir's controversial work with U.S. law enforcement*. The Logic.
<https://thelogic.co/news/exclusive/ottawas-ai-ethics-test-didnt-consider-palantirs-controversial-work-with-u-s-law-enforcement/>
- Hill, K. (2020, January 18). The Secretive Company That Might End Privacy as We Know It. *The New York Times*. <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- Hinton, G. E., & Frey, B. J. (1995). Using neural networks to monitor for rare failures. *Proceedings of the 37th Mechanical Working and Steel Processing Conference*.
- How Canada's unique research culture has aided artificial intelligence. (2017, April 11). *The Economist*. <https://www.economist.com/the-americas/2017/11/04/how-canadas-unique-research-culture-has-aided-artificial-intelligence>
- Hull, M. S. (2012). Documents and Bureaucracy. *Annual Review of Anthropology*, 41(1), 251–267. <https://doi.org/10.1146/annurev.anthro.012809.104953>
- Hunt, R. (2018). *The Heart's Content: Media and Marketing after the Attention Economy*. Concordia University.
- Hunt, R., & McKelvey, F. (2019). Algorithmic Regulation in Media and Cultural Policy: A Framework to Evaluate Barriers to Accountability. *Journal of Information Policy*, 9, 307–335. JSTOR. <https://doi.org/10.5325/jinfopoli.9.2019.0307>
- Immigration, R. et C. C. (2022, January 24). *Avis—L'analyse avancée des données pour aider les agents d'IRCC à trier les demandes de visa de résident temporaire et en faciliter le traitement*. Immigration et citoyenneté. <https://www.canada.ca/fr/immigration-refugies-citoyennete/nouvelles/avis/analyse-aider-demandes-visa-resident.html>

Immigration, Refugees and Citizenship Canada. (2021, October 7). *COW - Third-Party Vendors*. <https://www.canada.ca/en/immigration-refugees-citizenship/corporate/transparency/committees/cow-jun-10-2021/third-party-vendors.html>

Industry, Science and Technology Canada. (1989). *Networks of Centres of Excellence: Report of the International Peer Review Committee and report of the Minister's Advisory Committee*. Government of Canada. <https://publications.gc.ca/site/eng/9.871662/publication.html>

Innovation, Science and Economic Development Canada. (2021). *Minister Champagne presents Global Partnership on Artificial Intelligence priorities for 2021* [News Release]. ISED. <https://www.canada.ca/en/innovation-science-economic-development/news/2021/06/minister-champagne-presents-global-partnership-on-artificial-intelligence-priorities-for-2021.html>

Ipsos. (2022). *Global Opinions and Expectations about AI*. <https://www.ipsos.com/sites/default/files/ct/news/documents/2022-01/Global-opinions-and-expectations-about-AI-2022.pdf>

Jeffrey, B. (2010). *Divided loyalties: The Liberal Party of Canada, 1984-2008*. University of Toronto Press.

Jobin, A., Ienca, M., & Vayena, E. (2019a). Artificial Intelligence: The global landscape of ethics guidelines. *arXiv:1906.11668 [Cs]*. <http://arxiv.org/abs/1906.11668>

Jobin, A., Ienca, M., & Vayena, E. (2019b). *Artificial Intelligence: The Global Landscape of Ethics Guidelines*.

- Jobin, A., Ienca, M., & Vayena, E. (2019c). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Jones, M. (2022). Towards Civil Strategization of AI in Germany. An investigation of the involvement of civil society in the making of the German National Artificial Intelligence Strategy. *HIIG Discussion Paper Series*, 2022(1), 21. <https://doi.org/10.5281/ZENODO.6091638>
- Jones, M., Wester, M., & McKelvey, F. (2023). *Public Inputs as Training Data: Consultation, Calibration and Machine Learning Political Orders [Unpublished Manuscript]*.
- Karlin, M. (2017, July 5). Responsible AI in the Government of Canada: A Sneak Peek. *Code for Canada*. <https://medium.com/code-for-canada/responsible-ai-in-the-government-of-canada-a-sneak-peek-973727477bdf>
- Karlin, M. (2018, March 18). A Canadian Algorithmic Impact Assessment. *Supergovernance*. <https://medium.com/@supergovernance/a-canadian-algorithmic-impact-assessment-128a2b2e7f85>
- Karlin, M. (2022, March 18). [Interview by N. Gertler].
- Keung, N. (2023, June 2). *How AI is helping Canada keep some people out of the country. And why there are those who say it's a problem*. Toronto Star. https://www.thestar.com/news/canada/how-ai-is-helping-canada-keep-some-people-out-of-the-country-and-why-there/article_ea4dc715-debf-50e1-94c2-fda48a91dd09.html
- Köstler, L., & Ossewaarde, R. (2022). The making of AI society: AI futures frames in German political and media discourses. *AI & SOCIETY*, 37(1), 249–263. <https://doi.org/10.1007/s00146-021-01161-9>

- Koulish, R., & Evans, K. (2021). *Punishing with Impunity: The Legacy of Risk Classification in Immigration Detention* (SSRN Scholarly Paper 4022546).
<https://papers.ssrn.com/abstract=4022546>
- Kuchinskaya, O. (2014). *The Politics of Invisibility: Public Knowledge about Radiation Health Effects after Chernobyl*. MIT Press.
- Kuziemski, M., & Misuraca, G. (2020). AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. *Telecommunications Policy*, 44(6), 101976. <https://doi.org/10.1016/j.telpol.2020.101976>
- La Déclaration de Montréal IA responsable*. (2018). Déclaration de Montréal IA responsable.
<https://declarationmontreal-iaresponsable.com/la-declaration/>
- Lang, K. J., Waibel, A. H., & Hinton, G. E. (1990). A time-delay neural network architecture for isolated word recognition. *Neural Networks*, 3(1), 23–43.
[https://doi.org/10.1016/0893-6080\(90\)90044-L](https://doi.org/10.1016/0893-6080(90)90044-L)
- Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press.
- Lea, T. (2020a). *Wild policy: Indigeneity and the unruly logics of intervention*. Stanford University Press.
- Lea, T. (2020b). *Wild Policy: Indigeneity and the Unruly Logics of Intervention*. Stanford University Press. <https://doi.org/10.1515/9781503612679>
- Lepage-Richer, T., & McKelvey, F. (2022). States of computing: On government organization and artificial intelligence in Canada. *Big Data & Society*, 9(2), 205395172211233. <https://doi.org/10.1177/20539517221123304>
- Liboiron, M. (2021). *Pollution Is Colonialism*. Duke University Press.

- Light, B., Burgess, J., & Duguay, S. (2018). The walkthrough method: An approach to the study of apps. *New Media & Society*, 20(3), 881–900.
<https://doi.org/10.1177/1461444816675438>
- Linder, T., Jones, M., & McKelvey, F. (2023, January 12). Opinion | Toronto police consultation on AI lacks sufficient public engagement. *Toronto Star*.
<https://www.thestar.com/opinion/contributors/2023/01/12/toronto-police-consultation-on-ai-lacks-sufficient-public-engagement.html>
- Macdonald, D. (1985a). *Report of the Royal Commission on the Economic Union and Development Prospects for Canada: Volume 1*. Privy Council Office.
- Macdonald, D. (1985b). *Report of the Royal Commission on the Economic Union and Development Prospects for Canada: Volume 2*. Privy Council Office.
- MacKinnon, P. K. (1987). *An overview of AI and recent related government activities*. Ministry of State for Science and Technology.
<https://publications.gc.ca/site/eng/9.895747/publication.html>
- Margetts, H. (2022). Rethinking AI for Good Governance. *Daedalus*, 151(2), 360–371.
https://doi.org/10.1162/daed_a_01922
- Marres, N. (2020). For a situational analytics: An interpretative methodology for the study of situations in computational settings. *Big Data & Society*, 7(2), 2053951720949571.
<https://doi.org/10.1177/2053951720949571>
- MaRS. (n.d.). *What We Do*. MaRS Discovery District. Retrieved November 30, 2023, from <https://www.marsdd.com/what-we-do/>
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed*

- processing, volume 1: Explorations in the microstructure of cognition: Foundations*. MIT Press. <https://direct.mit.edu/books/monograph/4424/chapter/189409/The-Appeal-of-Parallel-Distributed-Processing>
- McIntyre, C. (2018, July 3). Discovering MaRS: The untold story of Canada's largest and most controversial innovation hub. *The Logic*.
- McKelvey, F. (2018, May 21). *Use the Charter to guide AI governance*. Policy Options. <http://policyoptions.irpp.org/magazines/may-2018/use-the-charter-to-guide-ai-governance/>
- McKelvey, F., Rajabiun, R., & McPhail, B. (2022, February 2). *AI accountability can't be left to the CRTC*. Policy Options. <https://policyoptions.irpp.org/magazines/february-2022/ai-accountability-crtc-oversight/>
- McStay, A. (2018). *Emotional AI: The Rise of Empathic Media*. SAGE Publications Ltd.
- Metz, C. (2019, March 27). Turing Award Won by 3 Pioneers in Artificial Intelligence. *The New York Times*. <https://www.nytimes.com/2019/03/27/technology/turing-award-ai.html>
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679. <https://doi.org/10.1177/2053951716679679>
- Morozov, E. (2013). *To Save Everything, Click Here: The Folly of Technological Solutionism*. PublicAffairs.
- Moss, E., Watkins, E. A., Singh, R., Elish, M. C., & Metcalf, J. (2021). *Assembling Accountability: Algorithmic Impact Assessment for the Public Interest*. Data & Society Research Institute.

- Neff, G., & Stark, D. (2004). *Permanently Beta: Responsive Organization in the Internet Era* (P. N. Howard & S. Jones, Eds.; pp. 173–188). SAGE Publications.
- Neural Computation & Adaptive Perception Program*. (2005, December 1). CIAR.
<https://web.archive.org/web/20051201124703/http://www.ciar.ca/web/home.nsf/pages/ncap>
- New York State Attorney General. (2023, May 10). *Attorney General James Secures \$615,000 from Companies that Supplied Fake Comments to Influence FCC’s Repeal of Net Neutrality Rules*. <https://ag.ny.gov/press-release/2023/attorney-general-james-secures-615000-companies-supplied-fake-comments-influence>
- Paltiel, G. (2021). The political imaginary of National AI Strategies. *AI & SOCIETY*.
<https://doi.org/10.1007/s00146-021-01258-1>
- Pangborn, M. H. (2010). *Defamiliarization*.
<https://doi.org/10.1002/9781444337839.wbelctv1d001>
- Parviainen, J. (2020). “*We’re Flying the Plane While We’re Building It*”: Epistemic Humility and Non-Knowledge in Political Decision-Making on COVID-19. 9, 6–10.
- Paul, K. T., Vanderslott, S., & Gross, M. (2022). Institutionalised ignorance in policy and regulation. *Science as Culture*, 31(4), 419–432.
<https://doi.org/10.1080/09505431.2022.2143343>
- Paul, R. (2022). Can critical policy studies outsmart AI? Research agenda on artificial intelligence technologies and public policy. *Critical Policy Studies*, 16(4), 497–509.
<https://doi.org/10.1080/19460171.2022.2123018>
- Peretti, J. (2017, July 30). Palantir: The ‘special ops’ tech giant that wields as much real-world power as Google. *The Guardian*.

- <https://www.theguardian.com/world/2017/jul/30/palantir-peter-thiel-cia-data-crime-police>
- Pestre, D. (2008). Challenges for the Democratic Management of Technoscience: Governance, Participation and the Political Today. *Science as Culture*, 17(2), 101–119.
<https://doi.org/10.1080/09505430802062869>
- Peters, J. D. (2013). Calendar, Clock, Tower. In J. Stolow (Ed.), *Deus in Machina: Religion, Technology, and the Things in Between* (1st ed, pp. 25–42). Fordham University Press.
- Pringle, R. (2020, February 1). *Controversial Clearview AI app could “end privacy.” So, what now?* CBC. <https://www.cbc.ca/news/science/clearview-app-privacy-1.5447420>
- PSPC. (2021). *Procurement: Standing Committee on Public Accounts—Procurement modernization*. <https://www.tpsgc-pwgsc.gc.ca/trans/documentinfo-briefingmaterial/pacp/2021-05-27/p2-eng.html>
- Public Safety Canada. (2022, September 20). *Combatting online terrorist and violent extremist content* [Backgrounders]. <https://www.canada.ca/en/public-safety-canada/news/2022/09/combating-online-terrorist-and-violent-extremist-content.html>
- Pylyshyn, Z. W. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences*, 3(1), 111–132.
<https://doi.org/10.1017/S0140525X00002053>
- Pylyshyn, Z. W. (1986). *Computation and cognition: Toward a foundation for cognitive science*. MIT Press. <https://doi.org/10.7551/mitpress/2004.001.0001>
- Radu, R. (2021). Steering the governance of artificial intelligence: National strategies in perspective. *Policy and Society*, 40(2), 178–193.
<https://doi.org/10.1080/14494035.2021.1929728>

- Reisman, D., Schultz, J., Crawford, K., & Whitaker, M. (2018). *Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability*. AI Now Institute.
<https://ainowinstitute.org/wp-content/uploads/2023/04/aiareport2018.pdf>
- Roberge, J., Morin, K., & Senneville, M. (2019). Deep Learning's Governmentality. The Other Black Box. In A. Sudmann (Ed.), *The Democratization of Artificial Intelligence: Net Politics in the Era of Learning Algorithms (Edition 1)* (pp. 123–143). transcript Verlag.
<https://doi.org/10.14361/9783839447192>
- Robertson, K., Khoo, C., & Song, Y. (2020). *To Surveil and Predict: A Human Rights Analysis of Algorithmic Policing in Canada*. Citizen Lab, University of Toronto.
<https://citizenlab.ca/2020/09/to-surveil-and-predict-a-human-rights-analysis-of-algorithmic-policing-in-canada/>
- Salée, D. (2010). Penser l'aménagement de la diversité ethnoculturelle au Québec: Mythes, limites et possibles de l'interculturalisme. *Politique et Sociétés*, 29(1), 145–180.
<https://doi.org/10.7202/039959ar>
- Sanders, N. E., & Schneier, B. (2023, January 15). Opinion | How ChatGPT Hijacks Democracy. *The New York Times*. <https://www.nytimes.com/2023/01/15/opinion/ai-chatgpt-lobbying-democracy.html>
- Schué, R. (2021, November 17). *Comment des étudiants étrangers francophones sont refusés par Ottawa*. Radio-Canada; Radio-Canada.ca. <https://ici.radio-canada.ca/nouvelle/1842994/etudiants-afrique-canada-immigration-universites-quebec-cegep-refus-permis>

- Senneville-Robert, M. (2021). *Reconfiguration des liens de collaboration entre acteurs industriels et universitaires de la recherche en intelligence artificielle à Montréal et à Toronto*. Université du Québec, Institut national de la recherche scientifique.
- Silcoff, S. (2020, December 18). *Element AI sold for \$230-million as founders saw value mostly wiped out, document reveals*. The Globe and Mail.
<https://www.theglobeandmail.com/business/article-element-ai-sold-for-230-million-as-founders-saw-value-wiped-out/>
- Simondon, G. (2012). *Du mode d'existence des objets techniques*. Aubier.
- Sloane, M., Moss, E., Awomolo, O., & Forlano, L. (2020). *Participation is not a Design Fix for Machine Learning*. 7.
- Smuha, N. A. (2021). From a 'race to AI' to a 'race to AI regulation': Regulatory competition for artificial intelligence. *Law, Innovation and Technology*, 13(1), 57–84.
<https://doi.org/10.1080/17579961.2021.1898300>
- Sold a Lie*. (2022, September 29). CBC. <https://gem.cbc.ca/the-fifth-estate>
- SOSCIP. (n.d.). *Mission & Core Values*. SOSCIP. Retrieved November 30, 2023, from <https://www.soscip.org/mission-diversity-core-values/>
- Standing Committee on Access to Information, Privacy and Ethics. (2022). *Facial Recognition Technology and the Growing Power of Artificial Intelligence*.
- Stark, L. (2019). Facial recognition is the plutonium of AI. *XRDS: Crossroads, The ACM Magazine for Students*, 25(3), 50–55. <https://doi.org/10.1145/3313129>
- Stark, L., & Hutson, J. (2021). *Physiognomic Artificial Intelligence* (SSRN Scholarly Paper ID 3927300). Social Science Research Network. <https://papers.ssrn.com/abstract=3927300>

- Sun, S., Zhai, Y., Shen, B., & Chen, Y. (2020). Newspaper coverage of artificial intelligence: A perspective of emerging technologies. *Telematics and Informatics*, 53, 101433.
<https://doi.org/10.1016/j.tele.2020.101433>
- Taeihagh, A. (2021). Governance of artificial intelligence. *Policy and Society*, 40(2), 137–157.
<https://doi.org/10.1080/14494035.2021.1928377>
- The end of anonymity? Facial recognition app used by police raises serious concerns, say privacy advocates | CBC Radio. (2020, January 22). In CBC.
<https://www.cbc.ca/radio/thecurrent/the-current-for-jan-21-2020-1.5434328/the-end-of-anonymity-facial-recognition-app-used-by-police-raises-serious-concerns-say-privacy-advocates-1.5435278>
- Thorpe, C., & Gregory, J. (2010). Producing the Post-Fordist Public: The Political Economy of Public Engagement with Science. *Science as Culture*, 19(3), 273–301.
<https://doi.org/10.1080/09505430903194504>
- Treasury Board of Canada (Director). (2021). *AI procurement for a digital world—Artificial Intelligence: Ethics and Responsibility Built-in* [Video].
<https://www.youtube.com/watch?v=LmyDFed5ZnI>
- Tsing, A. L. (2015). *The mushroom at the end of the world: On the possibility of life in capitalist ruins*. Princeton University Press.
- Tunney, C. (2019, November 5). *RCMP launches review of its social media monitoring operation* | CBC News. CBC. <https://www.cbc.ca/news/politics/rcmp-social-media-review-1.5346741>

- Ulnicane, I., Knight, W., Leach, T., Stahl, B. C., & Wanjiku, W.-G. (2021). Framing governance for a contested emerging technology: insights from AI policy. *Policy and Society*, 40(2), 158–177. <https://doi.org/10.1080/14494035.2020.1855800>
- Université de Montréal. (2018). *The Declaration—Déclaration de Montréal IA responsable*. <https://montrealdeclaration-responsiblaei.com/the-declaration/>
- Venturini, T. (2010). Diving in magma: How to explore controversies with actor-network theory. *Public Understanding of Science*, 19(3), 258–273. <https://doi.org/10.1177/0963662509102694>
- Voirol, O. (2005). Les luttes pour la visibilité. Esquisse d’une problématique. *Réseaux*, 129–130(1–2), 89–121.
- Webber, V., MacDonald, M., Duguay, S., & McKelvey, F. (2023). Pornhub and Policy: Examining the Erasure of Pornography Workers in Canadian Platform Governance. *Canadian Journal of Communication*, 48(2), 381–404. <https://doi.org/10.3138/cjc.2022-0044>
- Wehling, U. B., Peter. (2011). The politics of non-knowing: An emerging area of social and political conflict in reflexive modernity. In *The Politics of Knowledge*. Routledge.
- Wester, M. (2023). *Procurement Circuit under Machine Learning Political Order: Governance of, through, and for AI* [Masters, Concordia University]. <https://spectrum.library.concordia.ca/id/eprint/992057/>
- Wylie, B. (2023, December 24). *ISED’s Bill C-27 + AIDA. Part 10: The Origins of AIDA*. <https://biancawylie.com/iseds-bill-c-27-aida-part-10-the-origins-of-aida/>
- Zarsky, T. (2016). The Trouble with Algorithmic Decisions: An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making. *Science*,

Technology, & Human Values, 41(1), 118–132.

<https://doi.org/10.1177/0162243915605575>