

Workshopping AI: Who's at the Table?

Elia Rasky¹

ABSTRACT

In 2017, the Canadian federal government launched the “Pan-Canadian Strategy on Artificial Intelligence,” an ambitious plan to make Canada “a global leader in AI.” As part of this plan, the government sought to stimulate discussion about the ethical and societal implications of AI by sponsoring a series of AI & Society workshops. Hosted by the Canadian Institute for Advanced Research (CIFAR), these workshops brought together academics, engineers, and policymakers to discuss the impact of AI on healthcare, education, the modern workplace, Indigenous communities, and other areas. In its reports, CIFAR describes the AI & Society workshops as inclusive, diverse forums that allow actors from a range of different disciplinary, occupational, and ethnic backgrounds to express their opinions and concerns about AI. This paper investigates whether the AI & Society workshops are truly inclusive, or whether they privilege the voices and perspectives of some actors over others. It will be argued that, by inviting only “experts,” “thought leaders,” and “community leaders” to participate, the workshops systematically exclude laypeople and average consumers of technology. This is highly problematic since average consumers bear many of the social costs of advancements in AI. After critiquing the workshops, the paper proposes ways to amplify the voices of regular users of AI in public and intellectual discourse.

KEYWORDS: artificial intelligence, public forums, participatory democracy, deliberative democracy

RÉSUMÉ

En 2017, le gouvernement fédéral canadien a lancé sa « stratégie pancanadienne en matière d'intelligence artificielle (IA) », un plan ambitieux qui vise à faire du Canada « un chef de file en IA ». Dans le cadre de cette stratégie, le gouvernement a cherché à stimuler la discussion sur les implications éthiques et sociétales de l'IA, en parrainant une série d'ateliers sur le thème IA & société. Organisés par l'Institut canadien de recherche avancée (CIFAR), ces ateliers ont rassemblé des personnes du milieu académique, des ingénieurs.e.s ainsi que des décideur.e.s afin de discuter des impacts de l'IA sur la santé, sur l'éducation, sur le monde du travail, sur les communautés autochtones, et sur d'autres milieux.

Dans ses rapports, le CIFAR décrit les ateliers IA & Société comme des forums inclusifs et diversifiés, qui permettent aux parties prenantes de différentes disciplines, milieux professionnels et origines ethniques d'exprimer leurs opinions et leurs inquiétudes sur l'IA. Cet article vise à vérifier l'inclusivité de ces ateliers, cherchant à savoir si certaines voix ne sont pas privilégiées par rapport à d'autres. Il sera avancé qu'en invitant seulement les "expert.e.s", les "leaders d'opinion" et les "chef.fe.s de file communautaires" à participer à ces forums, ces derniers excluent systématiquement les profanes et les utilisateur.rice.s commun.e.s de ces technologies. Cela est hautement problématique en ce que ce sont principalement ces dernier.e.s qui ressentent les impacts sociaux de l'IA. Après une critique de ces ateliers, cet article présente des pistes de solution pour amplifier les voix des utilisateur.rice.s commun.e.s au sein des discours publics et intellectuels sur l'IA et ses impacts.

MOTS-CLÉS : intelligence artificielle, forums publics, démocratie participative, démocratie délibérative

¹ PhD Candidate, Department of Politics, York University.

Introduction

Over the past twenty years, artificial intelligence (AI) has become an integral part of the economy and society, impacting the lives of most people in a direct or indirect way. Beginning as a simple technology capable of solving math problems in a laboratory, AI now performs a range of complex functions in the financial sector, healthcare sector, public administration, transportation, law, and policing. Governments use AI to alert police of criminal activity online, and to determine whether individuals are eligible for welfare benefits, bail, and parole. In business, AI is used to make hiring decisions, assess the credit risk of loan applicants, and target advertising to consumers.¹ Healthcare providers rely on sophisticated algorithms to diagnose illnesses and predict the future medical needs of patients.² Of course, it is not just powerful public and private institutions that use AI. Average citizens and consumers have frequent interactions with AI systems, whether they conduct web searches, translate speech into text, or ask personal AI assistants for music or restaurant recommendations.

As AI has been widely adopted in the public and private sectors, and has become part of the lives of ordinary people, concerns have been raised about the possible negative social and economic impact of this technology. One of the main concerns, expressed by workers, politicians, and economists, is that the proliferation of AI could lead to mass unemployment, as work currently performed by humans is taken over by AI systems. Experts suggest that jobs in transportation, data operations, legal services, and customer service are likely to disappear in the near future due to advancements in AI.³ Another concern is that personal AI assistants and other AI-based tools may act against the interests of their users by, for example, harvesting personal information without user consent or providing users with biased recommendations. As Aguirre et al. explain: "A user who searches for a product using an AI system might assume that the results that are returned are the most relevant, highest quality, or best value, but in fact, such systems often prioritize results that provide the most financial benefit to the software designers, and algorithmic news feeds may prioritize user engagement time and advertising dollars over users' desire for true and useful information."⁴ Without laws or rules in place to regulate AI assistants, these high-tech tools are free to deceive, manipulate, and grift the users they are supposed to serve.

Computer scientists and software engineers warn that AI systems can behave in unpredictable ways, performing actions that are totally unintended by

¹ Kristin Johnson, "Automating the Risk of Bias," *George Washington Law Review* 87, no. 5 (2019), 1214 at p1215.

² Wendy Glauser, "AI in Health Care: Improving Outcomes or Threatening Equity?" *The Canadian Medical Association Journal* (2019) 21.

³ Judith Clifton, Amy Glasmeier & Mia Gray, "When Machines Think for Us: The Consequences for Work and Place," *Cambridge Journal of Regions, Economy and Society* 13, no. 3 (2020) 5.

⁴ Anthony Aguirre, Gaia Dempsey, Harry Surden & Peter B. Reiner, "AI Loyalty: A New Paradigm for Aligning Stakeholder Interests," *IEEE Transactions on Technology and Society* 1, no. 5 (2020) 128.

their users and designers. Although these actions are often harmless, they can sometimes have devastating consequences: An AI-driven truck may choose to collide with pedestrians in order to get to its destination faster.⁵ Sometimes, AI systems are not responsive to certain types of users—for example, women or people of colour—or engage in outright discriminatory behaviour.⁶ This can be seen with AI systems that assist companies with hiring employees. These “hiring bots” are designed to assess the educational qualifications and work experience of job applicants by scanning their resumes for certain key words. After conducting these assessments, the hiring bots identify the strongest applicants, and recommend them to hiring managers. What researchers have discovered, however, is that hiring bots tend to assess white male applicants more favourably than other types of applicants due to biases in their training data.⁷ This has disturbing implications: if hiring bots become widely adopted, and their biases are not eliminated, then they will greatly reinforce social and economic inequities. The potential dangers of AI to society go far beyond issues of unpredictability, bias, and discrimination. The risks are particularly acute in authoritarian states. As many political scientists have pointed out, AI systems “are showing their potential for abetting repressive regimes and upending the relationship between citizen and state, thereby accelerating a global resurgence of authoritarianism.”⁸ In China, Zimbabwe, and other countries with autocratic regimes, AI systems are being used to monitor communications among citizens, identify and silence dissidents, and speed up the mobilization of police and military forces against protestors.

As we can clearly see, the rapid development and deployment of AI raises many concerns with regard to economic security and equality, consumer welfare, human safety, and civic rights. To address these concerns, governments must harness the knowledge and resources of corporations, the nonprofit sector, the academic community, and the general population. People from different sectors of society must work together to ensure that new AI technologies serve the needs of society, and do not violate basic ethical and moral principles. Recognizing that good AI governance requires collaboration, governments have established forums where public, private, and nonprofit actors can discuss AI-related issues and develop regulatory policies and ethical frameworks for AI. These forums have been created not only by national governments, but also by international bodies such as the United Nations and the European Commission.

⁵ Alicia Solow-Niederman, “Administering Artificial Intelligence,” *Southern California Law Review* 93 (2020) 633 at 636.

⁶ James A. Rodger & Parag C. Pendharkar, “A Field Study of the Impact of Gender and Users’ Technical Experience on the Performance of Voice-Activated Medical Tracking Application,” *International Journal of Human-Computer Studies* 60, no. 5 (2004), 529 at 538; Ayanna Howard & Jason Borenstein, “The Ugly Truth About Ourselves and Our Robot Creations: The Problem of Bias and Social Inequality,” *Science and Engineering Ethics* 24 (2018) 1521 at 1525.

⁷ Johnson, *supra* note 1 at 1223.

⁸ Steven Feldstein, “The Road to Digital Unfreedom: How Artificial Intelligence is Reshaping Repression,” *Journal of Democracy* 30, no. 1 (2019) 40 at 41.

Typically, only individuals with certain levels of expertise, authority, or influence are allowed to participate in these forums, although this is not always the case.

In 2017, the Canadian federal government launched the “Pan-Canadian Strategy on AI,” an ambitious plan to “sustain and advance Canada’s leadership in AI research and innovation.”⁹ As part of this plan, the government provided funding for public forums that would explore the ethical implications and societal repercussions of AI. Hosted by the Canadian Institute for Advanced Research (CIFAR), these forums brought together individuals from a range of different disciplinary and occupational backgrounds, including academic researchers, civil servants, policymakers, and social activists. The government and CIFAR hoped that, by breaking down the barriers between disciplinary and occupational groups, and fostering dialogue between them, the forums could “holistically address the challenges and opportunities that AI presents.”¹⁰ In its reports, CIFAR discusses the importance of bringing multiple perspectives to bear on AI-related problems, arguing that these problems cannot be solved by single actors or organizations acting alone. The institute also identifies “inclusion” and “diversity” as the key organizing principles of the AI forums.

The aim of this paper is to critically analyze the AI forums hosted by CIFAR, investigating whether they truly allow for a diversity of voices and perspectives. The central argument of the paper is that the CIFAR forums on AI are elite-dominated spaces that only provide a platform to experts and professionals from the private, public, and nonprofit sectors. Lay people—that is, people with no advanced scientific and technical knowledge or political power—are excluded from these forums, even though they bear most of the social and economic costs of advancements in AI. The paper is divided into three sections. The first section provides background information on the CIFAR forums on AI, situating these forums within a larger global context. The second section discusses the benefits of including lay people in scientific debates, especially debates around AI policy and the role of AI in society. This section also shows how the CIFAR forums are completely devoid of lay or non-expert voices. The third section proposes ways to meaningfully include lay people in public debates about AI, drawing upon theories of deliberative and participatory democracy. The third section is followed by a conclusion. Although the CIFAR forums on AI are ongoing, this paper only focuses on the forums that were conducted between 2018 and 2020.

Before proceeding to the first section, it is worth quickly discussing the methodology and conceptual framework on which this paper is based. It relies on both documentary analysis and theoretical reasoning to build and present its

⁹ Canadian Institute for Advanced Research, “AICan 2019: Annual Report of the CIFAR Pan-Canadian AI Strategy,” *CIFAR* (2019) at 1

¹⁰ Canadian Institute For Advanced Research, “Rebooting Regulation: Exploring the Future of AI Policy in Canada,” *CIFAR* (2019) at 10.

arguments. All information about the CIFAR forums was derived from annual reports, workshop guides, and workshop briefs and summaries published by CIFAR during the years 2019 and 2020. Although these documents lack detailed information about forum participants, they clearly illustrate the exclusive and elite nature of the forums. It is quite evident from these documents that lay people were not invited to participate, and that their views were not solicited or sought after in any way by the forum organizers. This is highly problematic since lay people are especially vulnerable to the social and economic disruptions caused by AI. Not only this, but lay people possess valuable knowledge about AI that experts simply do not have. This knowledge, which is drawn from personal experience and ground-level observation, is highly relevant for the design and implementation of AI-related policies. In the future, governments must sponsor citizen forums on AI where lay people can study, critique, propose, and shape AI-related policies in a meaningful way. For these forums to be successful, they must be based on the principles of deliberative democracy. These principles will be expounded in the final section of this paper.

1. The Emergence of AI Forums in Canada and Abroad

All over the world, governments and policymakers have been grappling with the social, economic, and political challenges posed by AI. There is a desire among governments to protect citizens from the detrimental effects of AI discussed above. Considerable amounts of time, money, and energy have been spent by governments on the creation of laws that guide the development of AI and impose conditions on its use. One government that has been particularly committed to regulating AI has been the European Union. In April 2016, the European Parliament introduced the General Data Protection Regulation (GDPR), which grants citizens the right to know how their personal data is being used by digital service providers and their AI systems.¹¹ Under the GDPR, digital service providers must clearly explain to their users, if asked, how their AI systems make calculations and decisions. A few years after the passage of the GDPR, the European Commission published the “White Paper on Artificial Intelligence,” which proposed a comprehensive regulatory framework for AI.¹² This framework included new rules regarding the selection of training data, the disclosure of information to users, and the deployment of “high-risk” AI systems. Around the same time, the Organization for Economic Cooperation and Development (OECD) established a set of principles on the development and use of AI. These principles were designed to ensure that all new AI technologies are compatible

¹¹ Bryce Goodman & Seth Flaxman, “European Union Regulations on Algorithmic Decision Making and a ‘Right to Explanation,’” *Association for the Advancement of Artificial Intelligence: AI Magazine* (2017 fall edition) 50 at 55.

¹² Erik Brattberg, Raluca Csematoni & Vanesa Rugova, “Europe and AI: Leading, Lagging Behind, or Carving Its Own Way?” *Carnegie Endowment for International Peace* (2020) at 29.

with human rights and democratic values.¹³ Following in the footsteps of the OECD, a number of countries developed their own principles on AI ethics, including Britain,¹⁴ Canada,¹⁵ and China.¹⁶

National and international governments have responded to the risks and threats posed by AI by introducing new laws, regulations, and guiding principles, all of which are intended to better align AI systems with the interests and values of society. Realizing the potential of AI to cause serious societal harm, governments have carried out initiatives to make AI systems safer, fairer, and more equitable. Alongside this effort to harmonize AI with social interests, there has been an attempt by governments to generate public debate and dialogue about the relationship between AI and society. Civil servants, academics, corporate personnel, activists, and average citizens are being invited to participate in government-sponsored "AI forums," where they can engage in deep, wide-ranging discussions about the ethical and societal implications of AI. One of the most notable public forums on AI is the "European AI Alliance," which was launched by the European Commission in 2018.¹⁷ This forum allows experts on AI to receive direct feedback from members of civil society and concerned citizens about new policies on AI ethics. Another important forum is the "AI for Good Global Summit," created in 2017 by the United Nations and the XPRIZE Foundation. The Global Summit brings together stakeholders from the public, private, and nonprofit sectors to discover ways that AI can be used for the benefit of humanity and the environment.¹⁸ In the United States, the Federal Accountability Office hosted an AI forum in Washington, D.C., simply called the "Forum on Artificial Intelligence." The participants in this forum were all experts from major research institutions and regulatory agencies, such as OpenAI, Microsoft Research Lab, the National Institute of Standards and Technology, and the Financial Industry Regulatory Authority.¹⁹ For two days, these experts discussed the impact of AI on cybersecurity, transportation, criminal justice, and financial services, with each participant delivering a presentation on his or her area of expertise.

Given the popularity of AI forums around the world, it is no surprise that such forums have been established in Canada, a country with a long history of creating

¹³ Denise Carter, "Regulation and Ethics in Artificial Intelligence and Machine Learning Technologies: Where Are We Now? Who is Responsible? Can the Information Professional Play a Role?" *Business Information Review* 37, no. 2 (2020) 60 at 62.

¹⁴ The Committee on Standards in Public Life, "Artificial Intelligence and Public Standards," (2020) at 6.

¹⁵ Fenwick McKelvey & Margaret MacDonald, "Artificial Intelligence Policy Innovations at the Canadian Federal Government," *Canadian Journal of Communication Policy* 44 (2019) 43 at 45.

¹⁶ Huw Roberts, Josh Cows, Jessica Morley, Mariarosaria Taddeo, Vincent Wang & Luciano Floridi, "The Chinese Approach to Artificial Intelligence: An Analysis of Policy, Ethics, and Regulation," *AI & Society* 36 (2021) 59 at 68.

¹⁷ Brattberg, Csernatonni & Rugova, *supra* note 12 at 26.

¹⁸ James Butcher & Irakli Beridze, "What is the State of Artificial Intelligence Governance Globally?" *The RUSI Journal* 164 (2019) 88 at 93.

¹⁹ United States Government Accountability Office, "Artificial Intelligence: Emerging Opportunities, Challenges, and Implications," *Report to the Committee on Science, Space and Technology* (2018) at 52.

public forums on important public policy topics.²⁰ In 2017, the Canadian federal government launched the “Pan-Canadian Strategy on AI,” a program which aims to make Canada a “global leader in AI.”²¹ With a budget of \$125 million, the Pan-Canadian Strategy on AI has led to the creation of new AI research and training facilities in Toronto and Montreal, and the expansion of Canada’s AI talent pool. Most importantly, this federal program has led to the emergence of public forums that explore the ethical and societal implications of AI. These forums are planned, organized, and run by the Canadian Institute for Advanced Research (CIFAR), an illustrious nonprofit research institution with close ties to the federal government and Canadian universities. There are two types of forums hosted by CIFAR: AI & Society workshops and AI Futures Policy Labs. Participants in these forums not only discuss the myriad problems raised by AI, but also develop policy responses to these problems.²²

CIFAR began organizing the AI & Society workshops on April 9, 2018, when it invited teams of researchers from around the world to submit proposals for workshops. By May 28, CIFAR had received 39 proposals from 21 different countries.²³ After reviewing the proposals, CIFAR started the process of planning and scheduling the workshops. Between March 2019 and January 2020, 10 AI & Society workshops were held in 8 different cities: Toronto, Montreal, San Francisco, Santa Monica, Honolulu, London, Coventry, and Paris. Each workshop looked at the relationship between AI and a particular social institution or segment of society—for example, the healthcare system, the electoral system, the arts, children, Indigenous communities, and the environment. The average number of participants per workshop was 25, and most participants were academics, civil servants, and activists. Under the guidance of a facilitator, participants developed policy recommendations to address some of the ethical issues around AI. Some of these recommendations included the creation of independent data trusts or data banks, the inclusion of societal stakeholders in the design of AI systems, and the promotion of “algorithmic literacy” among children and parents.²⁴

²⁰ Steve Patten, “Democratizing the Institutions of Policymaking: Democratic Consultation and Participatory Administration,” *Journal of Canadian Studies* 35, no. 4 (2001), 230; Gregory Inwood, *Continentalizing Canada: The Politics and Legacy of the Macdonald Royal Commission* (University of Toronto Press, 2005) at 97; Genevieve Fuji Johnson, “Deliberative Democratic Practices in Canada: An Analysis of Institutional Empowerment in Three Cases,” *Canadian Journal of Political Science* 42, no. 3 (2009) 679 at 683.

²¹ Elissa Strome, Naser Faruqi and Rémi Quirion, “Canada First to Adopt Strategy for Artificial Intelligence,” *UNESCO Natural Sciences Sector Press Release* (2018) 1.

²² Canadian Institute for Advanced Research, “Workshop Brief—Beyond Arms Control: Regulating Defence and Security AI Technologies,” CIFAR (2019) at p1; Canadian Institute for Advanced Research, “Workshop Brief—Sustainability in the Digital Ages,” CIFAR (2019) at p1.

²³ Canadian Institute, *supra* note 9 at 28.

²⁴ Canadian Institute for Advanced Research, “Workshop Brief—Generation AI: Establishing Global Standards for Children and AI,” CIFAR (2019) at 1; Canadian Institute for Advanced Research, “Workshop Brief—AI-Powered Information Ecosystems and Democracy,” CIFAR (2019) at 1.

Alongside the AI & Society workshops, CIFAR organized another series of forums called the AI Futures Policy Labs, which aimed to educate policymakers and “emerging policy leaders” about the potential risks and benefits of AI.²⁵ The Policy Labs took place in 2018 and were carried out in partnership with the Brookfield Institute for Innovation, a nonprofit organization that seeks to help Canadians transition to a knowledge-based economy. There were three different stages to the Policy Labs. During the first stage, participants were asked to sit through a presentation by an AI expert, who explained how AI systems function and make decisions, and how they are currently being used in different sectors of society. During the second stage, participants were divided into small groups of five or six, with each group being asked to think about and discuss a particular AI technology. The technologies chosen for this exercise included AI systems that perform legal analytics, calculate life insurance premiums, make hiring decisions, and carry out other important functions.²⁶ The groups were presented with the following discussion questions:

(1) What different individuals or groups in society (stakeholders) may be affected by this technology? (2) How might each stakeholder be affected? (3) If this technology continues to advance, what might it look like a year from now? (4) Are there different ways that the technology could evolve (both positive and negative)? (5) What would the future look like without this technology?²⁷

In addition to answering these questions, the groups were tasked with developing specific policy recommendations. During the third stage of the Policy Labs, all participants reconvened to share what they learned and express their general thoughts about AI.

In one of its 2019 annual reports, CIFAR explains how academics, researchers, civil servants, and activists all over the world are examining, exploring, and discussing AI-related issues in isolation from each other. Instead of coming together to share knowledge and ideas about AI, these actors remain confined within their own professional and disciplinary spaces:

Many discussions about AI tend to take place in siloed or non-inclusive environments and do not accommodate the broad range of perspectives and subject-matter experts that are necessary to holistically address the challenges and opportunities that AI presents. It is important that all groups and organizations in the AI community work to bridge gaps in engagement and understanding.²⁸

²⁵ Canadian Institute, *supra* note 10 at 3.

²⁶ Canadian Institute for Advanced Research, “AI Futures Policy Lab: Toronto Pilot Summary,” *CIFAR* (2019) at 4.

²⁷ Canadian Institute for Advanced Research, “AI Futures Policy Lab Facilitator Guide,” *CIFAR* (2020) at 6.

²⁸ Canada Institute, *supra* note 10 at 10.

CIFAR believes that the ethical and societal problems raised by AI can only be solved if the barriers between academics, civil servants, and other actors are removed. Inter-disciplinary and inter-sectoral dialogue is essential for the creation of effective AI policies and regulations. By establishing the AI & Society workshops and AI Futures Policy Labs, CIFAR hoped to “bridge academic, civil service, nonprofit, and private sector perspectives [on AI].”²⁹ CIFAR imagined that these forums could serve as a kind of linchpin for the AI community, bringing different actors together to exchange and cross-pollinate ideas.

2. Lay People, Technology Policy, and the CIFAR Forums

So far in this paper, we have looked at AI forums from a purely functionalist perspective. AI forums have been seen as a tool used by governments to find answers to some of the ethical and normative questions around AI, and to assist with the regulation and governance of AI. It is important to view AI forums within a broader political context, however. In many ways, the emergence of AI forums in Canada and other parts of the world is emblematic of a general shift in politics, a shift away from hierarchical decision-making and towards collaborative or cooperative decision-making. Many political scientists have observed how, since the 1980s, policymaking around the world has become more open and flexible, as governments have allowed nonprofit organizations, activists, academics, and corporations to help shape public policy.³⁰ It has become common for governments to solicit the views of civil society actors on a whole range of public policy issues, from housing and immigration to child poverty and healthcare. George Frederickson, the former president of the American Society for Public Administration, attributes the rise of collaborative or democratic policymaking to what he calls the “disarticulation of the state”—that is, the inability of the state to address new public policy issues, many of which transcend spatio-juridical boundaries.³¹ A similar belief is held by Carey Doberstein, a Canadian political scientist, who views collaborative policymaking as “an institutional solution to complex governance problems, like homelessness, that cannot be resolved by any single level of government, ministry, or sector acting alone.”³²

Over the past few decades, collaborative policymaking has been used extensively in the area of science, technology, and environmental policy. Non-

²⁹ *Ibid.*

³⁰ Lisa Blomgren Bingham, Tina Nabatchi & Rosemary O’Leary, “The New Governance: Practices and Processes for Stakeholder and Citizen Participation in the Work of Government,” *Public Administration Review* 65, no. 5 (2005) 547.

³¹ George Frederickson, “The Repositioning of American Public Administration,” *Political Science and Politics* 23, no. 4 (1999), p702.

³² Carey Doberstein, *Building a Collaborative Advantage: Network Governance and Homelessness Policymaking in Canada* (UBC Press, 2016) at p20.

state actors have been called upon by governments to help create policies on genetics research, food safety, carbon emissions, pollution, land use, technology funding, and other techno-scientific and environmental issues.³³ The non-state actors who are asked to participate in the creation of techno-scientific policies are often “elites”—that is, individuals with high levels of social capital, influence, and professional achievement. These elites usually come from the corporate, nonprofit, and academic sectors, and are experts in particular fields or disciplines, such as engineering, finance, environmental science, and ethics. It would be wrong, however, to suggest that governments are only interested in speaking with elites. On occasion, governments have consulted with “lay people” about important scientific and technological developments.³⁴ Lay people are defined as individuals with little technical knowledge, policymaking experience, or political and economic power—in other words, average citizens. Governments seem to believe that both elites and lay people should contribute to the policymaking process.

On the surface, it seems strange that governments would ask lay people about their views on genetics research, information technology regulation, and other esoteric topics. After all, lay people appear to lack the knowledge needed to understand complex scientific and technological issues, never mind help create techno-scientific policies. Perhaps governments are engaging with lay people on matters of science and technology because they are under political pressure to do so. In his article titled *Citizen Forums Against Technocracy*, Tobin Craig argues that lay people today wish to be involved in decisions about science and technology, despite their lack of expertise on these subjects:

On issue after issue, from environmental regulation to compulsory vaccination, from food safety standards to hazardous waste disposal to telecommunication technology policy—every one of which on its face seems to challenge the competence of the non-expert—citizens, or at least the most active and vociferous among them, are refusing to simply defer to the claim of scientific expertise and are demanding more say in the determination of public policies.³⁵

³³ Mark Reed, “Stakeholder Participation for Environmental Management: A Literature Review,” *Biological Conservation* 141 (2008), 2417 at 2418; Patten, *supra* note 20 at 232; Anne Kerr, “Rights and Responsibilities in the New Genetics Era,” *Critical Social Policy* 23, no. 2 (2003) 208 at 217; Alan Irwin, “The Politics of Talk: Coming to Terms With the ‘New’ Scientific Governance,” *Social Studies of Science* 36, no. 2 (2006) 299 at 305; Patricia Fitzpatrick, A. John Sinclair and Bruce Mitchell, “Environmental Impact Assessment Under the Mackenzie Valley Resource Management Act: Deliberative Democracy in Canada’s North?” *Environmental Management* 42 (2008) 1 at 5; Adela Maciejewski Scheer & Corina Hoppner, “The Public Consultation to the UK Climate Change Act 2008: A Critical Analysis,” *Climate Policy* 10 (2010) 261 at 262.

³⁴ Simon Joss, “Public Participation in Science and Technology Policy and Decision-Making—Ephemeral Phenomenon or Lasting Change?” *Science and Public Policy* 26, no. 5 (1999) 290; Ida-Elisabeth Andersen & Birgit Jaeger, “Scenario Workshops and Consensus Conferences: Towards More Democratic Decision-Making,” *Science and Public Policy* 26, no. 5 (1999) 331 at 332; Ron Hagendijk & Alan Irwin, “Public Deliberation and Governance: Engaging With Science and Technology in Contemporary Europe,” *Minerva* 44 (2006) 167.

³⁵ Tobin Craig, “Citizen Forums Against Technocracy? The Challenge of Science to Democratic Decision Making,” *Perspectives on Political Science* 43 (2014) 31.

In the past, Craig says, lay people were completely shut out of the science and technology policy domain, which was a rarefied space dominated by “technocrats” (scientists, engineers, and professional policymakers). Without rigorous training and specialized knowledge, lay people felt unqualified to create, critique, or evaluate techno-scientific policies, even though these policies affected their lives in deep and profound ways. Now, lay people are asserting their right to shape and influence these policies, no longer resigning themselves to a role as passive observers. They are also challenging the statements and decisions made by scientists and policy experts, whom they see as untrustworthy, shortsighted, and reckless.³⁶ As governments have come under pressure from an increasingly restless, exuberant, and skeptical public, they have felt compelled to open the policymaking process to lay people.

As lay people have demanded a voice in the techno-scientific policymaking process, they have been supported by scholars specializing in the sociology of science. In their books and articles, these scholars have argued that lay people should be allowed to create techno-scientific policies, or assist in their creation, not only because they have a democratic right to do so, but also because they possess valuable scientific and technological knowledge. As various studies have shown, lay people know more about scientific and technological issues than is commonly believed.³⁷ For example, one study conducted in the 1990s shows that lay people understand many of the technical procedures, logistical challenges, and institutional norms associated with genetics research.³⁸ According to Michel Callon, a prominent science and technology theorist, lay people and experts have different forms of scientific knowledge that are of equal importance.³⁹ While the knowledge of experts is highly abstract and universal, the knowledge of lay people (lay knowledge) is “localized”—in other words, rooted in a particular geographical location, social setting, or community. Also, expert knowledge is derived from rigorous academic study and clinical experimentation, whereas lay knowledge is derived from personal experience and casual observation. Both forms of knowledge must be combined in order to solve complex scientific and technological problems. In a sense, one can see the relationship between lay and expert knowledge as dialectical and complementary.

³⁶ Alison Shaw, “‘It Just Goes Against the Grain’—Public Understandings of Genetically Modified (GM) Food in the UK,” *Institute of Physics Publishing* 11 (2002) 273 at 279; Darrin Durant, “Resistance to Nuclear Waste Disposal: Credentialed Experts, Public Opposition and Their Shared Lines of Critique,” *Scientia Canadensis* 3 (2007) 1.

³⁷ Darrin Durrant, “Accounting for Expertise: Wynne and the Autonomy of the Lay Public Actor,” *Public Understanding of Science* 17 (2008) 5.

³⁸ Anne Kerr, Sarah Cunningham-Burley & Amanda Amos, “The New Genetics and Health: Mobilizing Lay Expertise,” *Public Understanding of Science* 7 (1998) 41 at 45.

³⁹ Michel Callon, “The Role of Lay People in the Production and Dissemination of Scientific Knowledge,” *Science, Technology & Society* 4, no. 1 (1999) 81 at 85.

As many scholars have pointed out, lay people have a better understanding than experts of the societal impacts and implications of technology.⁴⁰ This is because lay people are fully exposed to the socio-economic disruptions caused by technological development. Whenever a new technology is introduced, lay people bear the social and economic consequences. For example, when corporations started using “customer service bots” based on AI, ordinary workers lost their jobs and livelihoods. While experts observe the societal effects of technology from a safe distance, lay people observe these effects up close. Due to their lack of resources and social capital, lay people often have difficulty adapting to technological change, which they see as an uncontrollable, tumultuous, and even destructive force. They know more about the depredations of technology than experts, who are sheltered from whatever societal harm technology may cause. One could argue that, since lay people are viscerally and materially impacted by technological change, they should be involved in the regulation of technology and the creation of techno-scientific policies.

As we saw in the introductory section of this paper, AI technology has become ubiquitous since the 1990s, raising a large number of ethical questions. Lay people must be allowed to answer these questions (or help answer them), because they are especially vulnerable to the disruptions caused by AI, giving them an insight and perspective that experts do not have. Some of the most important ethical questions related to AI include the following: (1) How can governments use AI to enhance public security without compromising civic rights? (2) How much data should AI assistants be allowed to gather from users? (3) Should labour-displacing AI systems be banned? (4) Given the risks involved, should AI be used in hiring, sentencing, and loan eligibility assessment? (5) Who should be held legally responsible if an AI tool engages in discrimination? (6) How much oversight should humans have over AI? Not only can lay people answer these questions, but they can also put forward new questions that policymakers have not yet considered. Drawing from their personal experiences and observations, lay people can identify problems and risks associated with AI that experts have overlooked. Once AI policies have been established, lay people can help the developers of AI systems comply with these policies. For example, lay people can inform AI developers as to whether their algorithms are easily understandable or interpretable (since the introduction of the GDPR by the European Parliament, “algorithmic interpretability” has become a legal requirement). They can also help AI developers explain the logic or reasoning behind their algorithms to regular users (another legal requirement under the GDPR).

⁴⁰ Gregor Durrenberger, Hans Kastenholtz & Jeannette Behringer, “Integrated Assessment Focus Groups: Bridging the Gap Between Science and Policy?” *Science and Public Policy* 26, no. 5 (1999) 341–342; Callon, *supra* note 39 at 87.

If lay people have an important role to play in the governance of AI, then they must be allowed to participate in the public forums that explore AI-related issues. As we have seen, national and international governments have conducted public forums on AI since the mid-2010s. Sometimes, these forums generate ideas that directly lead to new policies and regulations. Given their valuable technical insights, and their intimate knowledge of the societal impacts of AI, lay people deserve to be represented at these forums. Unfortunately, many AI forums deny lay people a seat at the table, only granting access to elites from the private, public, and nonprofit sectors. This was certainly the case with the AI & Society workshops and the AI Futures Policy Labs, the Canadian forums hosted by CIFAR. The organizers of these forums only invited “policy leaders,” “community leaders,” “thought leaders,” “activists,” “academic researchers,” “private sector practitioners,” and “innovators” to participate, completely excluding members of the general public.⁴¹ Although many participants knew little about AI, they all exercised considerable political and economic power.

CIFAR describes the AI & Society workshops and the AI Futures Policy Labs as truly diverse and inclusive spaces that allow people from all walks of life to learn about AI, express their thoughts about this technology, and help shape future AI policy. In its public statements, CIFAR emphasizes that these forums not only include people of different occupational and disciplinary backgrounds, but also people of “different ethnic and religious backgrounds, personal lived experiences, and career stages”.⁴² CIFAR’s claims of inclusion are misleading, however, since its forums exclude lay people and average users of AI technology. This is made very clear by the Workshop Briefs produced by CIFAR, which summarize the proceedings and outcomes of each AI & Society workshop. The Workshop Briefs show that, between March 2019 and January 2020, all participants in the workshops were academics, researchers, civil servants, and other actors with special technical knowledge and policymaking authority.⁴³ Also, the individuals who led and facilitated the Workshops were mostly researchers from prestigious institutions, such as Carnegie Mellon University, Microsoft Research Lab, and the UN Institute for Disarmament Research. These individuals recruited participants for the workshops from among their professional networks instead of from among the general population.

Like the AI & Society workshops, the AI Futures Policy Labs were entirely comprised of elites. This is not at all surprising, given that the Policy Labs explicitly targeted policymakers and “emerging policy leaders.” In one of its 2019 reports titled *Rebooting Regulation*, CIFAR states that the purpose of the Policy Labs was to educate policymakers in Canada about the policy implications of AI:

⁴¹ Canadian Institute, *supra* note 9 at 28; Canadian Institute, “Information Ecosystems,” *supra* note 24 at 1.

⁴² Canadian Institute, *supra* note 10 at 10.

⁴³ Canadian Institute, “Information Ecosystems,” *supra* note 24 at 1; Canadian Institute, “Children and AI,” (2019) at 1; Canadian Institute, “Beyond Arms Control,” (2019) at 1.

Many policymakers lack awareness of current AI capabilities and applications and their associated policy implications. For this reason, there is an increasing need to build capacity for thinking about emerging technologies among policymakers across all sectors to ensure AI is developed, implemented, and governed in ways that will align with public interest objectives...This series brought together over 125 policymakers from across Canada to learn about existing and potential AI capabilities and applications, explore the policy implications of AI, and develop policy responses.⁴⁴

The same report states that about half of the participants in the Policy Labs were government employees.⁴⁵ These included employees of the Canadian federal government, the provincial governments of Ontario, Alberta, Quebec, and British Columbia, and the municipal governments of Toronto, Edmonton, and Vancouver.⁴⁶ The other half of participants were members of nonprofit organizations, as well as public and private research institutions, such as the National Research Council and the Deloitte Future of Canada Centre. All participants were expected to take the ideas from the Policy Labs and apply them in their own professional practices or share them with their colleagues. It is important to note that CIFAR does not provide lists of attendees for either the AI & Society workshops or the AI Futures Policy Labs, nor does it identify all the organizations that were represented in these forums. Moreover, any information about the educational, occupational, and disciplinary backgrounds of participants is rather vague. The lack of precise information raises questions about the transparency of CIFAR as a semi-public organization. One could argue that CIFAR should be subject to more rigorous reporting standards, given that it receives federal money to deliver its programs. Even without precise information, however, it is abundantly clear that the CIFAR forums were elite-dominated spaces.

By completely shutting lay people out of the AI & Society workshops and AI Futures Policy Labs, CIFAR not only denied an opportunity for lay people to express their thoughts and feelings on AI, but also for researchers, civil servants, and other elites to gain valuable technical and sociological insights from this overlooked constituency. Had lay people been included in the forums, they could have provided the other attendees with a window into the lives of ordinary people, explaining how AI affects average workers, families, and communities. They also could have provided input into the policy development process, helping to ensure that policy recommendations are realistic and address the actual concerns of average citizens. There is a possibility, however, that even if lay

⁴⁴ Canadian Institute, *supra* note 10 at 3.

⁴⁵ *Ibid* at 5.

⁴⁶ Canadian Institute, *supra* note 9 at 29.

people were allowed to participate in the forums, their voices would have been minimized and their ideas would have been ignored. It is important to realize that AI forums, like all public forums on major policy topics, are discursive spaces that reward certain types of communication and argumentation. As Jennifer Dodge explains in her article titled *Tensions in Deliberative Practice*, actors in public forums are rewarded for presenting arguments that are clear, cogent, and couched in scientific evidence.⁴⁷ This is easy for elites, many of whom have received technical training and are familiar with abstract scientific concepts and language. Lay people, however, may have somewhat more difficulty constructing and presenting arguments that conform to standards of scientific rationalism. Instead of basing their arguments on scientific claims, theories, and facts, lay people are more likely to base them on personal experience, feeling, and emotion. If lay people cannot argue effectively, then their views may not be taken seriously.

In the 1990s and 2000s, when lay people started participating in public forums on scientific issues, some scholars investigated whether or not these forums are truly empowering for lay people. What these scholars discovered was that, in many cases, public forums do not allow lay people to exercise meaningful control or influence over scientific matters.⁴⁸ One of these scholars, Anne Kerr, examined public forums on genetics research in the UK.⁴⁹ These forums were hosted by major scientific institutes like the Royal Society, and brought together scientific and medical experts, medical patients diagnosed with genetic diseases, caregivers, and community members. Kerr found that medical patients, caregivers, and other lay people were extremely limited in what they could say during these forums. Specifically, they were only allowed to discuss minor technical issues like the phrasing of patient consent forms, and were prohibited from commenting on important policy issues like the future direction of genetics research. She also found that the policy recommendations arising out of these forums only reflected the views of the experts who participated. One may wonder why lay people were invited to participate in these forums at all, given that their views and opinions were not welcomed. One possible explanation is that, by bringing experts and lay people together, these forums sought to enhance public trust in the genetics research establishment. The logic here is that, if lay people directly interact with experts on genetics research, then they will come to understand their point of view and ultimately support their decisions. Some scholars have argued, in fact, that many public forums are nothing but public relations exercises designed to convince lay people to support certain policy decisions.⁵⁰

⁴⁷ Jennifer Dodge, "Tensions in Deliberative Practice: A View From Civil Society," *Policy Studies* 4, no. 4 (2010) 384 at 387.

⁴⁸ Alan Irwin, "Constructing the Scientific Citizen: Science and Democracy in the Biosciences," *Public Understanding of Science* 10 (2001) 1 at 13.

⁴⁹ Ann Kerr, *supra* note 33 at 215.

⁵⁰ Alan Irwin, *supra* note 33 at 306.

Clearly, lay people have many obstacles to overcome if they wish to shape or influence public policy on AI, or on any scientific and technological issue for that matter. In theory, governments have opened the policymaking process to lay people, allowing them to propose, critique, and evaluate techno-scientific policies. In practice, however, governments often exclude lay people from the public forums and consultations that develop techno-scientific policies, as we saw with the AI & Society workshops and AI Futures Policy Labs. Even when lay people are included in these forums and consultations, their views and opinions can often be overshadowed by those of experts, whose conceptual knowledge and rhetorical acumen give them an advantage in public deliberations. Also, lay people are sometimes prevented from expressing their views on broad policy issues due to rules, restrictions, and informal constraints imposed by experts. Despite these various impediments to participation, lay people want to exercise some control over techno-scientific policy. This raises an important question: How can we amplify the voices of lay people in public debates about science and technology policy, specifically in relation to AI? It is towards this question that we now turn.

3. Giving Lay People a Voice

As the new collaborative governance regime developed in the 1990s and 2000s, many scholars theorized that democracy itself was being transformed, and that the relationship between citizens and the state was being redefined.⁵¹ These scholars used the term “participatory democracy” to describe the new form of democracy that was taking shape. Unlike previous forms of democracy which concentrated power in the hands of elected representatives, participatory democracy disperses power among the general population, allowing citizens to make decisions about issues that affect their lives. Related to the concept of participatory democracy is the concept of “deliberative democracy,” which assumes that a healthy and functional democratic system requires active *deliberation* among elected representatives, and between representatives and the citizens they serve. According to this view, it is important that citizens not only vote for politicians every election, but also engage in extensive dialogue and debate about public policies and issues, sometimes as members of public forums or “mini publics.” Some proponents of deliberative democracy have proposed the creation of new institutions that might help bring citizens into the policymaking process. These institutions have many different names, such as “policy forums,” “policy committees,” “citizen forums,” “citizen juries,” “working groups,”

⁵¹ Archon Fung & Erik Olin Wright, “Deepening Democracy: Innovations in Empowered Participatory Governance,” *Politics & Society* 29, no. 1 (2001) 5; Patten, *supra* note 20 at 222; Jenny Cameron & Deanna Grant-Smith, “Building Citizens: Participatory Planning and a Transformative Politics of Difference,” *Urban Policy and Research* 23, no. 1 (2005) 21.

“consensus conferences,” and “roundtables.”⁵² Although these institutions differ structurally and organizationally, they all seek to give citizens a voice in the creation, design, and evaluation of public policy.

As we contemplate ways to help citizens shape AI policy, there may be lessons to be drawn from deliberative democracy theory. Perhaps some of the institutions proposed by deliberative democracy scholars could be applied to ensure that citizens monitor, assess, and direct the development of AI technologies. One institution that may be suitable for this purpose is the “integrated assessment focus group,” described by Durrenberger, Kastenholz, and Behringer.⁵³ This institution brings average citizens together to discuss a particular social issue or policy topic, develop a position on the issue or topic, and submit a recommendation to government officials. Unlike conventional forums, which privilege expert voices and place limits on citizen participation, these focus groups are designed to “give [lay] people as much control over the process as possible, and to reduce expert knowledge, jargon, and teaching attitudes to a minimum”.⁵⁴ Although experts may provide information to these focus groups, they cannot directly participate in them, as this would risk disempowering citizens. Another institution that is worth considering is the “citizen panel” which is a “collection of lay citizens akin to a jury but charged with deliberating on policy issues with a high technical content.”⁵⁵ During the sessions of a citizen panel, citizens formulate questions about a particular issue, and then call upon experts to answer these questions in person.⁵⁶ The citizens then weigh the facts and evidence offered by the experts, and seek a consensus about the issue in question.

In 2004, the government of British Columbia created its own citizen forum called the BC Citizens' Assembly, which was greatly admired by participatory and deliberative democracy scholars at the time. The purpose of this forum was to examine the electoral system of British Columbia, and to decide whether the province should adopt a new system. The forum was composed of 160 individuals from across the province who were selected through an entirely random method. The participants represented a cross section of the BC population who varied in age, ethnicity, and education level, and “were employed in a range of sectors, from restaurant owners, computer programmers, factory hands, and

⁵² Ank Michels, “Innovations in Democratic Governance: How Does Citizen Participation Contribute to a Better Democracy?” *International Review of Administrative Sciences* 77, no. 2 (2011) 275 at 276; Manuel Fischer & Philip Leifeld, “Policy Forums: Why Do They Exist and What Are They Used For?” *Policy Science* 48 (2015) 363 at 364.

⁵³ Durrenberger, Kastenholz & Behringer, *supra* note 40 at 341.

⁵⁴ *Ibid* at 343.

⁵⁵ David H Guston, “Evaluating the First U.S. Consensus Conference: The Impact of the Citizens' Panel on Telecommunications and the Future of Democracy,” *Science, Technology & Human Values* 24, no. 4 (1999) 451.

⁵⁶ Andersen & Jaeger, *supra* note 34 at 331.

accountants, to retail sales people, teachers, ranchers, and child-care workers."⁵⁷ For almost a whole year, this diverse group of participants came together to learn about different electoral systems around the world, debate the merits of these systems, and critically evaluate BC's system. In the spring of 2004, participants travelled around the province, asking fellow citizens about their views on BC's electoral system. After months of learning, listening, and debating, the participants decided that BC should switch from a "first past the post" electoral system to a "single transferable vote" (STV) system. Although the STV system was rejected in a provincial referendum in 2005, the BC Citizens' Assembly still stands as an example of a successful, innovative, and empowering citizen forum. Critics may argue that the rejection of the STV system was also an indictment of the BC Citizens' Assembly, or a sign that the Assembly did not truly reflect the values and views of the BC population. According to many scholars, however, the majority of BC residents voted against electoral reform simply because they did not understand the STV system, nor did they understand the purpose and function of the BC Citizens' Assembly.⁵⁸ This widespread ignorance among residents was not the fault of the Assembly, but rather that of the provincial government, which utterly failed to inform its constituents about the Assembly and its proposal during the lead up to the referendum. The more residents knew about STV and the BC Citizens' Assembly, the more they were inclined to support electoral reform. Perhaps the BC Citizens' Assembly could serve as a model for a future citizen forum on AI.

If lay people wish to shape or influence public policy on AI, then they must adopt the cognitive and communication conventions associated with the policymaking realm. To be more specific, they must learn how to translate their fears, hopes, and beliefs about AI into concrete political demands, and then articulate these demands in civic and political environments. They must also learn how to speak to, question, and challenge the various elites who currently control and dominate public discourse on AI and society. One way for lay people to acquire these skills is by participating in citizen forums. According to many deliberative democracy scholars, citizen forums are not just "discussion platforms" where lay people can express their views on various issues. They are also "pedagogical devices" that allow lay people to expand their knowledge, improve their communicative and deliberative abilities, and become better advocates for their own interests. Frank Fischer, one of the foremost experts on deliberative democracy, argues that the main function of citizen forums is to "assist citizens in their effort to better understand and communicate their own

⁵⁷ Mark E. Warren & Hilary Pearse, "Introduction: Democratic Renewal and Deliberative Democracy," in *Designing Deliberative Democracy: The British Columbia Citizens Assembly*, Mark E. Warren & Hilary Pearse (eds) (Cambridge University Press, 2009) 1 at 10.

⁵⁸ Fred Cutler, Richard Johnston, R. Kenneth Carty, Andre Blais & Patrick Fournier, "Deliberation, Information, and Trust: The British Columbia Citizens' Assembly as Agenda Setter," in *Designing Deliberative Democracy: The British Columbia Citizens Assembly*, Mark E. Warren & Hilary Pearse (eds) (Cambridge University Press, 2009) 166 at 172.

ideas and interests in the face of power and opposition."⁵⁹ By participating in these forums, Fischer says, citizens or lay people gain a clearer sense of where they stand on complex societal issues. They also acquire the intellectual and rhetorical tools needed to act effectively, and assert themselves, in the civic and political realm.

Before creating a citizen forum on AI, a government must decide whether or not to allow academics, policymakers, and other elites to participate in the forum. If elites are allowed to participate, then they may completely take over the forum, stripping lay people of their voice and relegating them to the margins. On the other hand, forums that integrate elites and lay people could result in a fruitful exchange of information and ideas about AI, leading to the emergence of effective AI ethics policies. Under the right conditions, elites can teach lay people about the mechanics of AI, update them about recent AI developments, and present them with different AI policy options. Conversely, lay people can provide elites with an intimate account of the social, economic, and cultural effects of AI, giving them a better understanding of the societal and ethical issues at stake. For elites and lay people to cooperate, however, both groups must have an equal opportunity to speak, introduce discussion topics, and set priorities. Not only this, but elites and lay people must see each other as equals; elites must recognize that lay people have valuable ideas to contribute, while lay people must recognize that elites share many of their values and concerns. Of course, even if elites and lay people have mutual respect, they are still unequal in terms of communicative ability and technical knowledge. Also, elites have more time to invest in AI forum activities than lay people because they are professionals with flexible work schedules. As Nancy Fraser and other scholars have pointed out, mutual recognition and respect between social groups do not guarantee an equal distribution of economic resources or political/institutional authority.⁶⁰ To ensure that an "even playing field" exists within AI citizen forums, it may be necessary to grant lay people special institutional or procedural powers. For example, lay people may be granted the ability to make final changes to reports.

In his article titled *Professional Expertise in a Deliberative Democracy*, Fischer explains how elites and lay people have different ways of seeing the world because they possess different epistemological frameworks.⁶¹ Elites (especially scientific experts) see the world through a conceptual and positivist lens, while lay people see the world through a personal subjective lens. Given their different worldviews, elites and lay people may have difficulty working together in a forum setting. If a forum is run and facilitated well, however, then it can allow elites and lay people to forge an intellectual, emotional, and spiritual bond. An example of

⁵⁹ Frank Fischer, "Professional Expertise in a Deliberative Democracy: Facilitating Participatory Inquiry," *The Good Society* 13, no. 1 (2004) 21 at 24.

⁶⁰ Nancy Fraser, "Rethinking Recognition," *New Left Review* 3, no. 3 (2000) 107 at 110.

⁶¹ Fischer, *supra* note 59 at 26.

such a forum is the Responsible Science and Public Engagement Workshop, hosted by Durham University in the UK. This forum brought together scientists and members of the lay public to explore issues around the environment. During the forum, participants not only discussed climate change, pollution, and other threats to the environment, but also hiked through a forest near the venue where the forum was taking place. As the participants walked along the forest trails, they were asked to reflect on the relationship between humanity and nature, and to imagine what the world would look like after a mass environmental catastrophe.⁶² After the nature hike, the participants felt a sense of common purpose and shared humanity, despite their different social and occupational statuses and levels of knowledge/power. They expressed a desire to “move forward together rather than [stay] divided into established bunkers of “science” versus the “public.”⁶³ Perhaps a similar exercise can be carried out in a citizen forum on AI. For example, participants (elites and lay people) can be asked to read a fictional story in which society has lost control of AI, or in which AI is being used for nefarious purposes. Alternatively, participants can imagine a society that has effectively harnessed AI to increase productivity and improve people's lives. Such an exercise would put elite and lay participants “on the same page,” closing the intellectual and emotional gulf between them.

Conclusion

The AI forums hosted by CIFAR can be seen as part of a global movement to address the ethical issues around AI through broad-based dialogue. These forums were intended to be inclusive, diverse, and multidisciplinary spaces where individuals from different occupational and ethnic backgrounds can discuss the problems with AI and develop solutions to these problems. This paper showed, however, that CIFAR's AI forums were elite-dominated spaces that excluded lay people or average citizens. The exclusion of lay people was problematic for two different, but related, reasons. Firstly, AI systems produce social and economic instability for lay people, who have difficulty adapting to these systems due to their lack of resources and social capital. Many lay people have jobs that are vulnerable to automation, like clerical and customer service jobs, and do not understand how AI systems gather and process information, making them especially susceptible to data harvesting and manipulation. If lay people are shut out of the AI policymaking process, then their concerns about AI cannot be heard by governments, and the threats they face from AI cannot be mitigated. Secondly, lay people have valuable ideas to contribute to the AI policymaking process, because they truly understand the societal and ethical implications of

⁶² Phillip Robinson, Phil Macnaghten, Sarah Banks, Janie Bickersteth, Angela Kennedy, Yvonne Richardson, Sue Shaw & Ingrid Sylvestre, “Responsible Scientists and a Citizens’ Panel: New Storylines for Creative Engagement Between Science and the Public,” *The Geographical Journal* 180, no. 1 (2014) 83 at 85.

⁶³ *Ibid* at 86.

AI. Having directly experienced the social and economic disruptions caused by AI, lay people can critique and evaluate existing ethics policies on AI, suggesting minor and major ways to improve or update them. They can also identify ethical issues that experts have overlooked and propose new policies to address these issues. One can argue that governments cannot develop effective AI ethics policies if they do not consult with lay people or include them as stakeholders in the policymaking process.

When one reads CIFAR's annual reports and Workshop Briefs, it is not immediately clear why the organization excluded lay people from its AI forums. One possible explanation is that CIFAR, or more precisely the teams that organized its forums, did not believe that lay people knew enough about AI to make a meaningful contribution. This explanation is not satisfactory, however, because many of the participants in the forums knew little about AI. A more likely explanation has to do with CIFAR's desire to influence elite political discourse. From the very beginning, CIFAR wanted its AI forums to generate policy ideas that would reverberate through the corridors of power, stimulating discussion and debate among elites in the public, private, and nonprofit sectors. As a result, CIFAR invited people who were embedded in professional communities and policymaking circles to participate in its forums. The hope was that, once the forums were over, these people would spread AI policy ideas throughout government departments, nonprofit organizations, and corporate offices. Since lay people have little political and economic power, they could not help CIFAR promulgate policy ideas, making them useless. Field research on the subject would channel questions about the inclusion or exclusion of lay people in AI forums to lead organizers at CIFAR. What role, if any, do lay people play in the creation of AI policies? What conflicts and challenges are anticipated if lay people were to participate? And for those who contributed to successful citizen forums like the BC Citizens' Assembly, how did their voices influence public policy?

Drawing upon deliberative democracy theory, this paper proposed ways in which lay people can be brought into the AI policymaking process. It was suggested that citizen forums provide the key to citizen empowerment in the area of AI, allowing lay people to learn about the mechanics of AI, express concerns about this technology, and propose and shape policies that protect their interests. It was also argued that these forums can only be successful if certain conditions are met. Firstly, lay people must have the authority to set the agenda, introduce discussion topics, summon and dismiss experts, assemble final reports, and issue recommendations. Secondly, if elites are allowed to participate in these forums, then they must respect lay people and not attempt to rule over them. Going even further, it is important that a spiritual and emotional connection exist between elites and lay people, as this allows for ongoing cooperation and a fruitful exchange of ideas. We should hope that, in the coming decades, citizen forums on AI start to receive funding from governments, not just in Canada but

around the world. Elites can no longer be allowed to monopolize the AI policymaking space, especially as AI technologies start to transform every social, cultural, economic, and political institution.