FINAL REPORT

Policy Recommendations for AI and Online Harms Governance in Canada

GEN(Z)AI:

Final Report

Policy Recommendations for AI and Online Harms Governance in Canada

Helen A. Hayes and Fergus Linley-Mota

With Contributions From:
Madeleine Case, Alexander Martin,
Julian Lam, and Nonso Morah

Centre for MEDIA, TECHNOLOGY and DEMOCRACY

DoT
Dialogue on Technology Project
SFU Morris J. Wosk Centre for Dialogue

Mila

Project Leads
Helen Hayes and
Fergus Linley-Mota

Contributors
Madeleine Case
Alexander Martin
Julian Lam
Nonso Morah

Operations Lead
Sequoia Kim

Designers
Brian Morgan
Ibrahim Rayintakath
Mathilde Robert

Translator
Diane Robert

Special Thanks
Isabelle Corriveau
Emma Frattasio
Anna Jahn
Sequoia Kim
Taylor Owen

Centre *for* MEDIA, TECHNOLOGY *and* DEMOCRACY

The Centre for Media, Technology and Democracy is an interdisciplinary research organization housed at the Max Bell School of Public Policy at McGill University. Its work is dedicated to understanding and responding to the social, political, and policy challenges posed by the evolving information ecosystem and digital technologies.

DoT
Dialogue on Technology Project
SFU Morris J. Wosk Centre for Dialogue

The Dialogue on Technology Project (DoT) is the SFU Morris J. Wosk Centre for Dialogue's flagship initiative on technology and artificial intelligence. Drawing on the Centre's deep expertise in design, facilitation, consensus building, and conflict engagement, DoT brings together diverse stakeholders—community members, researchers, policymakers, industry, and civil society—to better understand and shape how AI and related technologies are transforming our lives and societies.

This project was generously supported by our partners

CIFAR
RONALDS ROADBURG FOUNDATION
WALTONS TRUST

CONTENTS

The Halifax Forum, photo by Rafe Ayub

FOREWORD

When we set out to design *Gen(Z)AI: Canada's Youth Assembly on Artificial Intelligence*, we were animated by an urgent conviction: that young people, among the most exposed to and affected by AI systems, should not merely be subjects of this technology, but architects of its governance. What followed—across seven months, four cities, one hundred in-person participants, and thousands of virtually engaged young Canadians—exceeded our initial ambitions.

Citizens' assemblies are a well-established democratic tool, used in Canada and globally to translate the complexity of public life into actionable insights. Building a novel intervention in the space therefore felt like an exercise in both creativity and academic rigour. Our three-day forum structure, which blended virtual literacy-building with in-person deliberation and hybrid consensus, was designed to honour the realities of our participants—17-23 year-old Canadians who have come of age alongside AI. We collaborated closely with subject matter experts, civil society advocates, and our extraordinary project team to develop programming that was intellectually rigorous and participatory, while always centering participants' lived experience as a form of expertise in its own right. At the core of this methodology is the premise that legitimacy in governance flows from genuine participation, and that ordinary people, given adequate information and structured space for

deliberation, are capable of extraordinary collective judgement. Gen(Z)AI has demonstrated that this insight applies with particular force to questions of AI governance.

For those who work in deliberative democracy, the theoretical case for citizens' assemblies is well-rehearsed. What is harder to convey, and what we hope this report in some measure communicates, is how it feels to witness one work. To sit in a room as a participant who arrived uncertain of their own expertise begins, over the course of three days, to speak with authority; to watch consensus emerge through the honest and generous navigation of disagreement; to see young people leave the process with a transformed sense of their own civic agency: this is the promise of deliberative democracy fulfilled.

We offer this report to those shaping the future of AI in Canada: to policymakers tasked with designing guardrails for a rapidly evolving technology; to civil society organizations, researchers, and practitioners seeking more meaningful ways to engage the public; and to members of the public who rightly expect to have a voice in decisions that will shape their lives. The recommendations presented here are evidence-based, consensus-built, and grounded in the realities of how AI and online harms are actually experienced by young people. They are written in the participants' own words, drafted and ratified through the deliberative process described in this report, and presented here without alteration. The contextual analysis and policy framing that surrounds them is our own, offered to situate their recommendations within Canada's current legislative landscape and to identify the pathways through which participants' policy goals could be operationalized.

We are confident that the recommendations drafted by Gen(Z)AI participants represent something rarely achieved in technology governance: the democratic voice of those who have the most to gain or lose, speaking clearly and in their own name.

Sincerely,

Helen A. Hayes & Fergus Linley-Mota
Gen(Z)AI Project Leads

Acknowledgements

A project of this ambition became real in large part due to the hard work of our Youth Fellows team—Maddie Case, Alexander Martin, Julian Lam, and Nonso Morah—who acted as facilitators, interpreters, advocates, and co-creators throughout the process. They brought warmth to spaces that might otherwise have felt procedural, and rigour to conversations that might otherwise have remained vague. To work alongside them was one of the many privileges of this project.

We are equally grateful for the guidance and expertise shared with us by our Gen(Z)AI Advisory Board[1] and our partners: Mila and Make.org, whose complementary expertise in AI and large-scale digital engagement made the full arc of this project possible. We would particularly like to thank Taylor Owen and Anna Jahn for their support of our vision for the project, and Aftab Erfan for instilling her trust in us to lead with care and commitment.

Gen(Z)AI was made possible by the generous financial support of The Waltons Trust, The Ronald S. Roadburg Foundation, and CIFAR.

Project Team: Fergus Linley-Mota, Helen Hayes, Maddie Case, Nonso Morah, Alexander Martin, and Julian Lam.

Photo by Tishan Baldeo

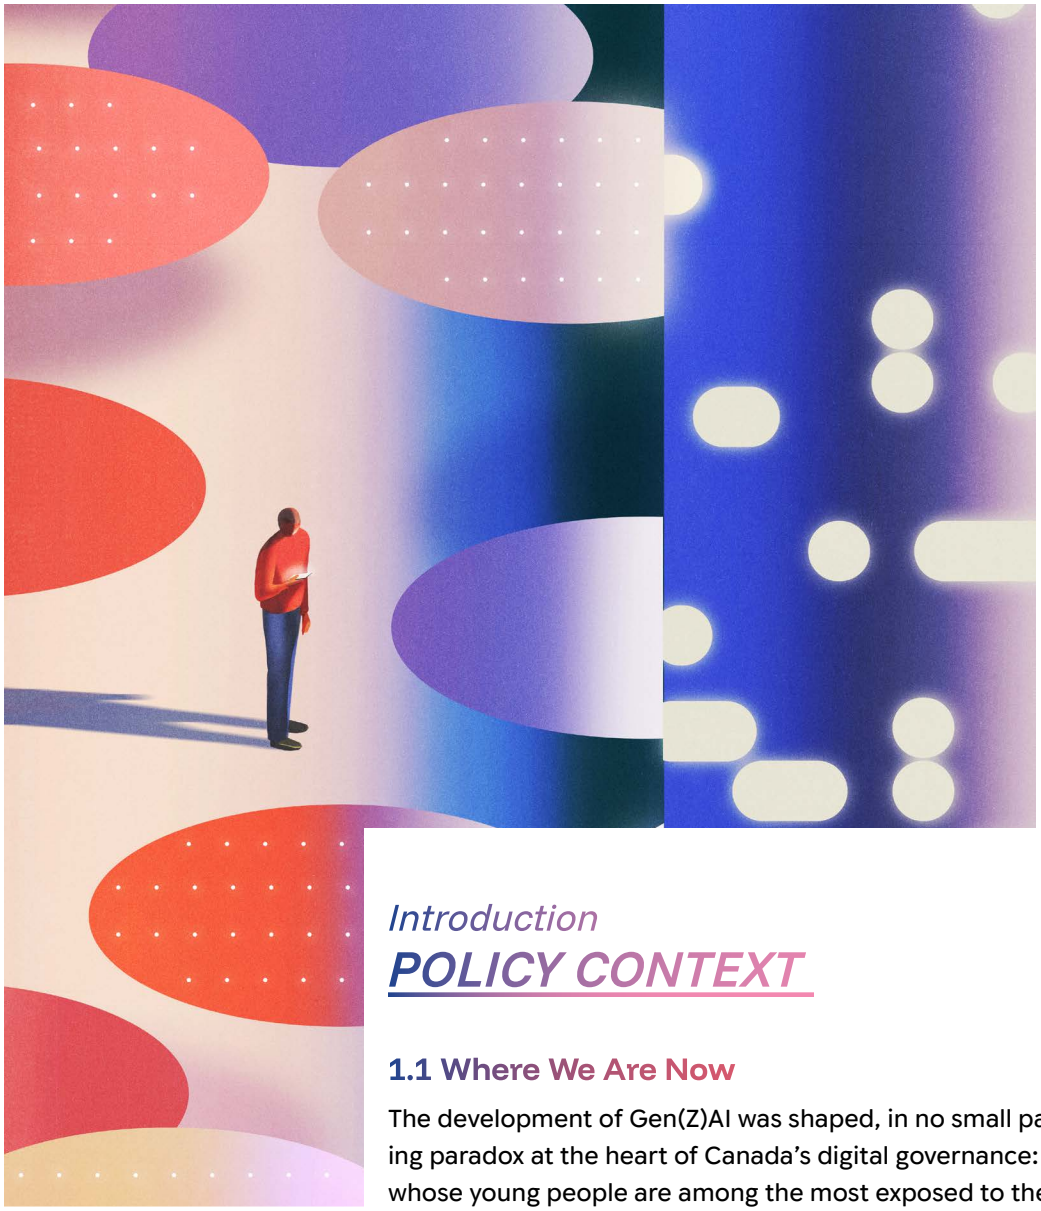

1

Introduction

POLICY CONTEXT

1.1 Where We Are Now

The development of Gen(Z)AI was shaped, in no small part, by a striking paradox at the heart of Canada's digital governance: that a country whose young people are among the most exposed to the harms of algorithmically-mediated life has, as of this writing, no binding legal framework to regulate either AI systems or online platforms. Understanding how Canada arrived at this moment, and why this represents both a crisis and an important opportunity, is essential context for the recommendations that this report advances.

In January 2025, when Parliament was prorogued following Prime Minister Justin Trudeau's announcement of his resignation, two landmark pieces of digital legislation died on the order paper. Bill C-63, *The Online Harms Act*, and Bill C-27, which paired federal privacy reform with the country's first dedicated AI regulatory statute, the *Artificial Intelligence and Data Act* (AIDA), failed to attain Royal Assent. AIDA's collapse, in particular, carried a specific weight of consequence: throughout its legislative cycle, the bill attracted sustained and substantive criticism from researchers, civil society organizations, and advocates who were concerned that it was too narrow in scope, too

The Montreal Forum, photo by Jack Devine

limited in enforcement capacity, and too inattentive to the concerns of affected communities. What remained were aging statutes designed for a pre-platform world, and a series of voluntary AI codes of conduct (for federal use cases only).

We write this report within the context of a government that has inherited this legislative deficit alongside an electorate that is increasingly aware of what is at stake if AI and online harms legislation is not prioritized and tabled. There is now broad cross-partisan recognition that reform cannot be deferred indefinitely, despite the government's explicit adoption agenda and investment in AI innovation. The question, then, is how Canada will act, and whether the framework that emerges will be proportional to the complexity of the problem, and genuinely responsive to those who bear its costs.

1.2 Fragmented Governance Landscape

Canada has never approached digital governance as a unified project. The responsibility for regulating the digital environment has long been distributed across institutions—the Office of the Privacy Commissioner, the Canadian Radio-Television and Telecommunications Commission, the Competition Bureau, the Justice Ministry, ISED, and the Department of Canadian Heritage—each operating under its own statutory mandate, with limited mechanisms for coordination and no shared architecture for coherent standards-setting.

This institutional fragmentation reflects and reinforces longer standing questions about how digital governance itself has been viewed and defined by the federal government. To date, online harms regulation has been conceived of primarily as a content problem concerning how platforms identify, limit, or remove material deemed harmful. Privacy reform has proceeded on a parallel track, focused on individual rights over personal data and the conditions of its lawful use. AI-specific regulation, as imagined under AIDA, proposed to address automated systems in commercial contexts while leaving recommender algorithms and affective AI systems largely untouched. And, competition policy has focused on market concentration and platform dominance, with its tools designed for an economic logic that predates the dynamics of the attention economy.

Each of these threads responds to a real and serious concern. What none of them does, either individually or in combination, is grapple with the digital environment as users actually experience it: as an integrated space in which AI systems, platform design, data practices, and the circulation of harmful content are not separate phenomena, but aspects of one online reality.

1.3 Governance Gaps

This structural mismatch between how governance is organized and how harm is experienced is the condition that Gen(Z)AI was designed to address. Existing AI safety debates, both in Canada and internationally, have been dominated by questions of frontier model capabilities and catastrophic risk. These are not trivial concerns, but their dominance has directed policymakers' attention toward speculative futures and away from AI systems already in operation. Online harms debates, meanwhile, have focused on content moderation, treating harm as a problem of visibility rather than of architecture.

The international community has made serious attempts to respond to these challenges. The European Union's *AI Act* and *Digital Services Act*, the United Kingdom's *Online Safety Act*, and Australia's *Online Safety Bill* each represent significant regulatory ambition. Yet each, in their own ways, continues to struggle with the intersection of AI and online harms as a unified problem, defaulting more often than not to the same siloed logic that has constrained Canadian efforts.

Young people experience no such division. For the participants of Gen(Z)AI, AI and online harms are not adjacent issues, but interconnected ones: the algorithmic systems that curate their social lives also amplify harmful content; the platforms that mediate their relationships are simultaneously the infrastructure through which AI-generated material circulates. The governance gap is therefore not simply a matter of missing legislation, but of a serious framing issue. Gen(Z)AI's focus on the intersection between online harms and AI legislation has sought to fill this gap.

2

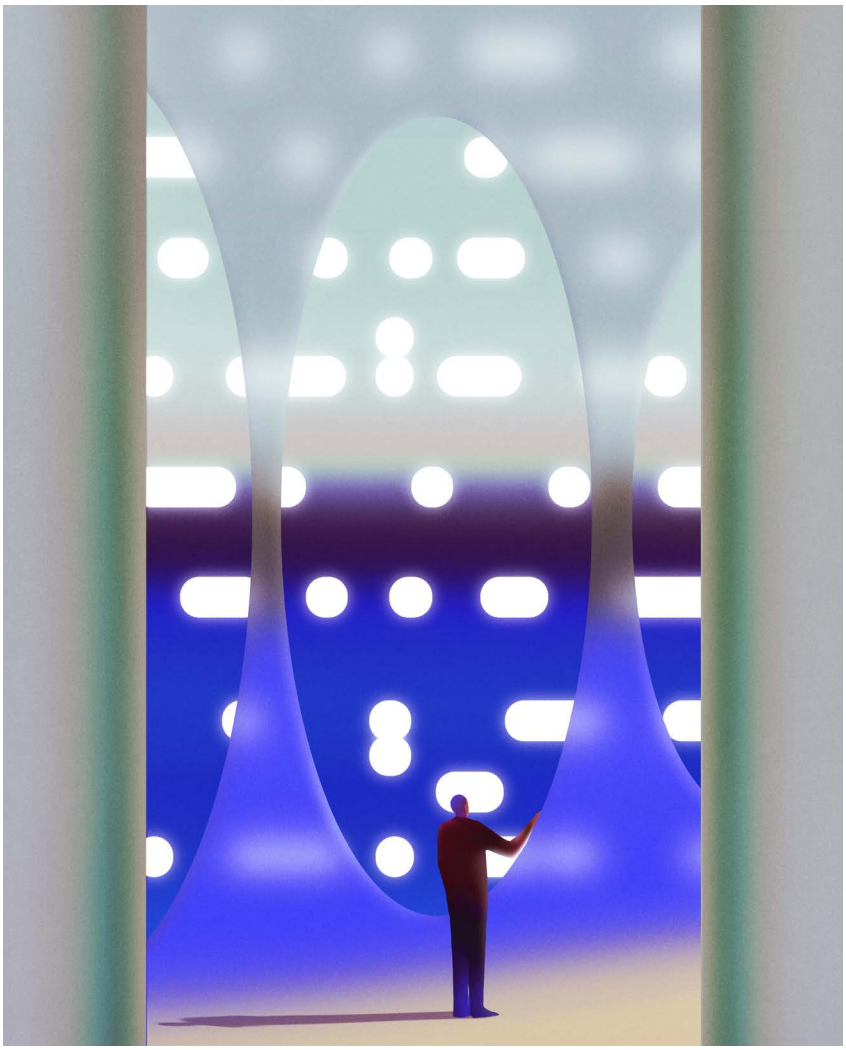

Methodology

A NOVEL YOUTH CITIZENS' ASSEMBLY MODEL

2.1 Design

Canada's governance landscape has failed young people in part because it has spoken *about* them rather than *with* them, routing their interests through expert proxies and industry consultations. Our methodological premise was that a different kind of governance process was both possible and necessary.

Gen(Z)AI engaged one hundred randomly selected young Canadians, aged seventeen to twenty-three, across four regional in-person deliberative forums held between November 2025 and March 2026. Each forum was convened in a different city and organized around a distinct policy domain identified as central to young people's digital lives: AI chatbots in Toronto, information integrity in Montreal, data privacy in

Vancouver, and age assurance in Halifax. The selection of these four domains map the terrain where AI systems and online harm converge most acutely in the everyday experience of young people, the precise intersection that existing governance frameworks have been least equipped to address.

Participants were recruited through a national outreach process and selected via civic lottery to reflect Canada's geographic, linguistic, and demographic diversity. Our civic lottery process,[2] a hallmark of citizens' assembly design, was central to the project's democratic logic and ensured that the assembly was constituted not by those already engaged with policy processes, but by ordinary young people whose experience of AI-mediated life qualified them as experts in the most relevant sense.

2.2 Process

Each regional forum in Toronto, Montreal, Vancouver, and Halifax followed a three-day structure adapted from the design principles of citizens' assemblies and calibrated to the practical and social realities of participants' lives. The process unfolded in three distinct but interconnected phases: 1) a virtual learning session that opened each forum; 2) two days of in-person structured deliberation; and, 3) a national digital engagement stage that extended the forum's reach to thousands of young Canadians beyond the in-person assembly. Prior to each virtual learning session, participants were provided with a policy resource pack that, without weighing in on the merits of different approaches, provided them with accessible information about the relevant policy landscape in both Canada and peer jurisdictions.

The Toronto Forum, photo by Tishan Baldeo

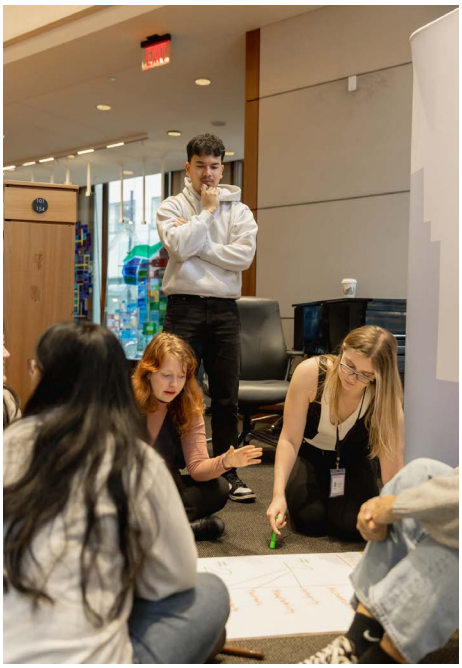

The Vancouver Forum, photos by Alison Boulier

2.2.1 Virtual Learning Day

Each forum opened with a virtual learning day, involving all one hundred Gen(Z)AI participants. This session was designed to establish a shared informational foundation, regardless of participants' prior familiarity with each forum's policy domain. A curated panel of subject-matter experts delivered short, accessible presentations and engaged directly with participants in a live question-and-answer session.[3] Experts were selected to represent a range of disciplinary and institutional perspectives and were briefed in advance on the importance of delivering evidence in a form that was accessible to non-specialist participants—a deliberate design choice that reflects one of the core commitments of our methodology. Expert presentations were structured to provide context on the relevant policy landscape, the specific harms and opportunities associated with the forum's topic, and the key questions that participants would be working through together in the following days.

2.2.2 In-Person Deliberation

The in-person deliberation, held over two days at host institutions in each regional city,[4] was the core of the Gen(Z)AI process. During this in-person session, participants moved through a structured sequence of activities in both small-group and plenary settings, with small-groups facilitated by our trained youth fellows[5] and plenary sessions led by the Gen(Z)AI project leadership.[6]

Facilitated Small Group Discussions: Participants were divided into small-groups of approximately five to seven people, each facilitated by a youth fellow. The small-group structure was designed to create the

conditions for genuine deliberation: intimacy, structured progression, and sufficient time to move from personal experience and values to the co-development of concrete policy positions. The youth fellows worked from a facilitation guide developed by the Gen(Z)AI co-leads, moving groups through successive stages of reflection, from individual perspective-sharing to collective recommendation-building. Groups were intentionally composed to reflect diversity of geographic origin, background, and prior experience with the forum's topic.

Values and Issues Mapping: a structured values and issues mapping exercise formed the bridge between the learning phase and the recommendation-building phase. Working first in their small-groups and then sharing back in plenary, participants identified the values they believed should underpin AI governance in the context of their forum's topic; surfaced the specific issues, tensions, and dilemmas that most concerned them; and mapped out where values and concerns came into conflict with one another. This exercise produced a structured picture of the landscape that directly informed the recommendation-building phase, ensuring that the policy proposals that ultimately emerged were anchored in participants' own articulated values.

Recommendation Building: using the values and issues map as a foundation, participants worked through a structured consensus-building process, supported by the youth fellows, to develop one to two concrete policy recommendations for each issue statement surfaced during the mapping phase. The aim was to produce recommendations that are specific, actionable, and grounded in the lived experience of young Canadians navigating AI-mediated digital life.

The Halifax Forum, photo by Rafe Ayub

Collective Reflection: participants reconvened in a "policy gallery walk" exercise to view and vote on their fellow participants' initial policy recommendations, and then, after synthesis of the recommendations presented during the gallery walk, reconvened in plenary to review, refine, and finalize those recommendations. In the spirit of collective decision-making, all recommendations were subject to a modified consensus or majority vote. This phase reinforced participants' sense of ownership over the forum's outputs and prepared them for their role as communicators and advocates for the recommendations they had collectively developed.

2.2.3 National Digital Engagement

Upon the completion of each regional forum, the policy recommendations produced through in-person deliberation were published on Make.org's civic dialogue platform: a bilingual, interactive tool designed to extend the deliberative reach of the process to young Canadians across the country. Beginning after the first forum and growing with each successive one, the platform supported three substantive forms of engagement, each designed to capture a different dimension of national youth opinion:

1. **Priority ranking**, through which users indicated which issues and recommendations they considered most urgent, producing a scaled and weighted picture of priorities that both complemented and validated the in-person forum outputs.
2. **Endorsement or challenge**, through which users indicated whether they supported, partially supported, or disagreed with a given recommendation, with space to explain their position in their own words. This mechanism was designed to surface the full range of opinion across the national youth population and ensure that minority or dissenting views were captured.
3. **Amendment**, through which users proposed refinements to the wording or scope of any recommendation, enabling a genuinely iterative process in which the broader youth community could contribute to and shape the work produced at each forum.

The results of this national digital engagement will be available in a companion report in May 2026.

2.3 A Methodology for Democratic Legitimacy

What distinguishes Gen(Z)AI from more conventional approaches to policy consultation is not simply its scale or its youth focus, but its underlying theory of democratic legitimacy. Standard consultation models, including written submissions, stakeholder roundtables, and public comment periods, can be procedurally compliant but are often substantively thin. This is because they aggregate inputs without contextual understanding and invite participation without creating the real conditions for it. Gen(Z)AI was designed on a different premise: that

deeper, deliberative forms of engagement—those that bring people into sustained, informed, and contextual exchange—can meaningfully strengthen the legitimacy of governance processes, particularly in domains where consultation risks becoming procedural or symbolic.

The "layered deliberation" architecture that structures the project—pairing intensive in-person assembly work with large-scale digital engagement—represented an attempt to hold two democratic values in productive tension: deep deliberation and broad participation. Deliberation, which requires time, facilitation, shared information, and the experience of genuine and contextual encounter with perspectives different from one's own, cannot, in our view, be scaled without loss. But deliberation that remains confined to a small group also carries its own legitimacy deficit. By sequencing the two (grounding the digital engagement phase in recommendations produced through rigorous in-person dialogue), Gen(Z)AI sought to extend democratic reach without sacrificing democratic depth.

This architecture also represents a direct challenge to what might be called the techno-solutionist reduction of democratic participation—or, more simply, the assumption that engagement can be adequately captured through clicks, digital rankings, and algorithmic aggregation. The in-person participants of Gen(Z)AI built understanding, navigated disagreement, and arrived at consensus through a process that required them to take genuine interest in one another, and to engage each other in real, shared presence. The recommendations that follow are the product of that process and it is precisely that provenance that gives them their claim to democratic authority.

The Halifax Forum, photos by Rafe Ayub

3

WHAT WE HEARD FROM YOUNG PEOPLE ACROSS CANADA

The Gen(Z)AI forums covered four distinct policy issues, each at the intersection of online harms and AI legislation: AI chatbots, information integrity, data privacy, and age assurance. This section aggregates what we heard from participants, organized by the cross-cutting themes that recurred across cities and topics. Forum-specific issue statements, produced by Gen(Z)AI participants, are reported in subsection 3.7.

3.1 Transparency

Across all four forums, one of the most consistent findings we heard was that existing transparency mechanisms, including privacy policies, consent notices, terms of service, and disclosure prompts, are systematically failing young users. In this context, participants first identified what could be called “disclosure fatigue,” a condition of indifference created by the sheer density of consent mechanisms that renders genuine engagement with them practically impossible. In our Vancouver Forum, participants characterized the terms and conditions governing

major AI platforms as functionally unreadable in everyday use. Several noted that services like ChatGPT feel designed to make meaningful engagement with underlying data practices unnecessary and invisible for users; consent, in this design logic, is registered at the moment of first use, before most users have any understanding of what they are tacitly consenting to by using the tool. The Vancouver forum surfaced perhaps the starkest expression of this dynamic: several participants reported having "given up on data privacy" as early as fourteen, assuming that extraction of their information had long outpaced any realistic capacity to contest it.

A second dimension of the transparency failure, named with particular force in the Montreal forum, concerned the automatic integration of generative AI tools in search. Examples such as Gemini's embedding in Google, they expressed, change the nature of conventional information retrieval, with meaningful implications for how users not only access but, importantly, assess the information they receive. In this sense, participants called for transparency not merely as a policy lever related to data practices, but also as a matter of basic legibility of digital tools in online spaces. They wanted to know, at the moment of their interaction, what kind of system they were engaging with and on what terms.

The third dimension of the transparency problem was articulated most clearly in the Toronto forum, but was echoed in all others: that AI systems are designed as black boxes, with limited visibility into how outputs are generated, how training data is constituted, or how user information is stored, reused, or incorporated into future AI model iterations. Participants expressed frustration about the opacity of AI systems and argued that this opacity seemed constitutive of the systems themselves. This was raised as a particular concern in the Halifax forum, where participants expressed uneasiness about the future use of sensitive data, including biometric data, collected and stored in the process of verifying users' ages. Transparency is not, in participants' views, a problem that better disclosure language will resolve alone. It is what participants described as a problem of systems design.

3.2 Individual Onus

Woven through, and deepening, the transparency problem is a structural condition that participants across all four forums named with remarkable consistency: the systematic displacement of responsibility for complex, system-level harms onto the individuals least equipped to bear it. Young people, participants noted, are expected to identify mis- and dis-information in saturated information environments, protect their own privacy against constant and largely invisible data extraction, self-regulate their engagement with AI systems deliberately designed to addict users, and navigate platforms whose inner workings remain deliberately opaque. Participants expressed the difficulty of being expected to do all of this without the institutional support, informational access, or structural leverage that would make it achievable.

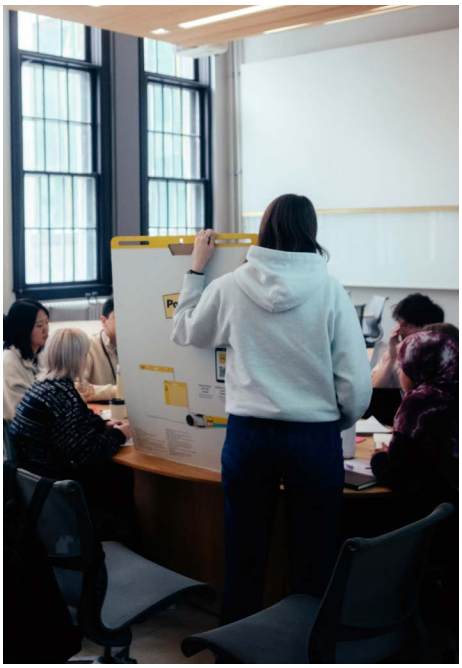

The Montreal Forum, photos by Jack Devine

In the Montreal forum, for example, participants pushed back against the framing of information integrity as a problem of individual media literacy. As one participant observed, the challenge is not whether any given piece of content is false, but rather that the sheer volume and velocity of AI-generated and algorithmically amplified material makes evaluation at scale structurally impossible for any individual user. For Montreal participants, consistent encounters with false or misleading information were a predictable output of platform environments that are optimized for engagement rather than informational quality, within which the burden of verification had been offloaded entirely onto users.

In the Toronto forum, the burden of individual responsibility manifested in its most personal form, as participants described the difficulty of disengaging from AI chatbot interactions, and the experience of noticing in themselves a reflexive reliance they had not chosen. The responsiveness and apparent emotional attunement of chatbot interfaces were understood by participants as design choices made to sustain engagement. To hold individuals responsible for navigating these dynamics, participants said, is to misplace responsibility in a way that is both practically unreasonable and ethically untenable. This sentiment was echoed in the Halifax forum, where participants noted concerns about age assurance measures displacing users from “mainstream” platforms and AI tools to potentially harmful online spaces, where safety burdens would fall disproportionately on individuals.

The Vancouver Forum, photo by Alison Boulier

3.3 Engagement Maximization

Across all four forums, participants were consistent in their views that many of the harms they described were outcomes of systems deliberately optimized for engagement. In Toronto, participants reflected at length on the role of addictive design in AI chatbots. They argued that the sycophancy of many chatbot systems is intended to sustain interaction, cultivate dependency, and maximize time-on-platform. The capacity to blur the boundary between human and machine relationships, to reinforce rather than interrogate a user's existing beliefs and emotional states, and to generate the false experience of being understood, are design choices made with the intention of driving commercial success. Several participants described their own experiences of cognitive offloading or emotional reliance that they found difficult to reverse, and linked these dynamics to design choices they had never consented to.

In Montreal, the same logic was expressed in participants' discussion of users' information environment. Recommendation algorithms, participants observed, are designed to surface what most reliably sustains and intensifies engagement, thus systematically favouring emotionally charged, sensational, and/or polarizing content. Participants noted that there was "no guarantee they would ever see the other side" of a contested issue within their algorithmically curated feeds because the incentive structure underlying content distribution assigns no commercial value to balance and considerable value to engagement. The proliferation of what

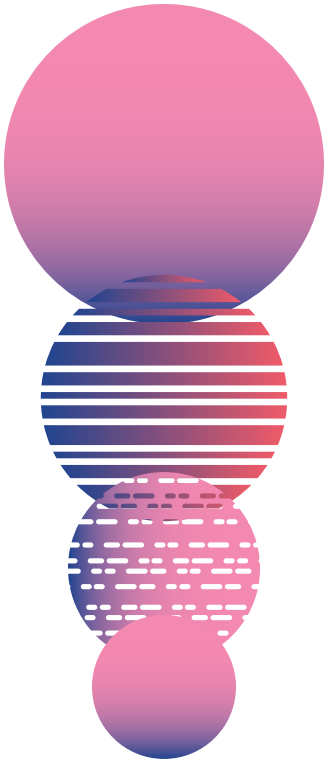

participants called “AI slop”—low-quality, mass-produced AI-generated content that circulates at scale—is itself a product of this logic.

In Vancouver, participants described both the opacity of data collection practices and the intentionality behind them—the sense that the systems collecting their data had been specifically designed to do so in ways that precede, surround, and outlast any individual act of consent. For many participants, the challenge, then, lay in addressing the misalignment between the commercial interests of platform companies and the wellbeing of the users whose data and attention those interests depend upon.

3.4 Accountability

Underlying the transparency failure, the individualisation of risk, and the design problem, participants consistently noted a lack of existing accountability for AI’s harms. Across all four forums, participants described a governance landscape in which the entities whose systems produce harm face no meaningful obligation to prevent it, report it, or answer for it; in which the internal mechanisms that do exist are difficult to access and produce no visible consequence; and in which the absence of enforceable external standards has abandoned the field to voluntary commitments that carry, in practice, no binding force.

This was articulated with particular urgency in the Toronto Forum, when participants described the ease with which AI chatbot systems can be prompted to produce harmful content, including material related to self-harm and suicidal ideation. Most AI chatbot platforms do contain internal reporting mechanisms, but participants described them as difficult to locate and unreliable in response. Participants noted that even when an issue is successfully reported, there is no assurance that it has been received in good faith, no visible consequence, and no indication that the system has been modified as a result.

The Montreal Forum surfaced accountability gaps of a related but distinct kind. Participants identified the absence of effective regulatory oversight of recommendation systems and the engagement-incentive structures driving them. This absence, they noted, allows platforms to profit from dynamics whose harms are well-documented and widely understood, without any obligation to address them. For participants, systems optimized for engagement are doing so at an unacceptable cost to the users they engage.

3.5 Generational Stakes

While participants in Gen(Z)AI spoke from their own experience, they were rarely only speaking for themselves. Across all four forums, participants articulated a concern that extended beyond their present circumstances to encompass those younger than themselves, including younger siblings, friends and future generations. This intergenerational awareness was an animating value throughout the process.

The Gen(Z)AI cohort, made up of participants aged seventeen to twenty-three, occupies a distinct position in the social history of AI technologies. They are the first generation to have grown up fully immersed in AI-mediated platforms, and the first to have had algorithmic recommendation systems shaping their information environment from an early age. Several Toronto participants described what it has meant to navigate relationships using chatbots as "friendly ears," and expressed concern for younger people who may begin to view chatbot systems as being more reliable than their real-life friends. The Vancouver Forum made the temporal dimensions of this concern concrete in a different way: the permanence of data. Participants were troubled by the fact that data being collected about them was not bounded or recoverable. Several questioned whether AI systems could meaningfully "forget" content once it had been provided, and expressed concern that data could remain embedded in model training datasets indefinitely and beyond an individual's awareness or control. Participants were particularly concerned about what this meant for data exposure in childhood, and the fact that harms incurred before young people had the legal standing or practical capacity to contest them could not simply be undone.

What gives these stakes their genuinely generational character is the compounding nature of the harms involved. A young person whose cognitive development is shaped by AI systems designed to cultivate dependency, or whose data is commodified before they have either the legal standing or the practical capacity to contest it, carries those experiences forward in ways that are not fully visible at the time they occur, and that governance frameworks oriented toward discrete, documentable harms are poorly equipped to address. Participants emphasized the need for policymakers to hold this longer view in order to understand that the governance decisions made now are not only about the harms already documented, but about the kind of people, relationships, citizens, and societies that our current trajectory may produce.

The Halifax Forum, photos by Rafe Ayub

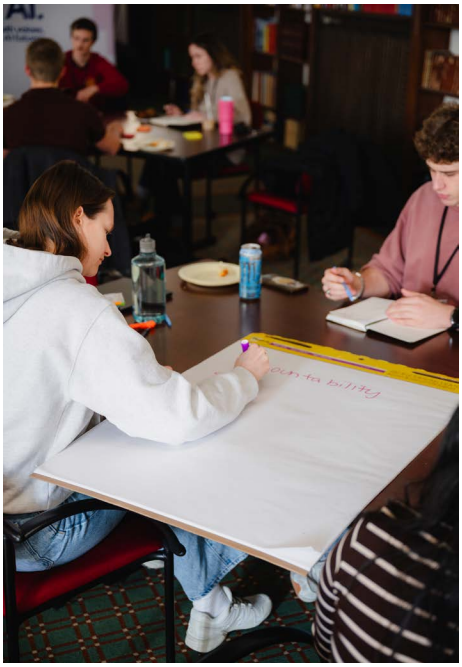

The Toronto Forum, photo by Tishan Baldeo

3.6 The Participation Deficit

Participants consistently described their general exclusion from the governance processes that shape their digital lives, and considered most youth engagement that has occurred to be tokenistic rather than substantive. This was particularly salient in discussions of age assurance, where the vulnerability of children and young people is routinely invoked as a justification for regulatory intervention, yet young people themselves remain largely absent from the decision-making spaces where those interventions are designed. Participants identified a pattern of reactive engagement: input solicited only after foundational decisions have already been made, constraining young people's role to reactive feedback, rather than proactive consultation.

The deeper argument that participants made asserted that governance of AI systems can only be adequate if it is grounded in the lived realities of those most profoundly affected by them. They were clear that young people are and should not merely be subjects of AI governance; that they are, by virtue of their experience, among its most qualified informants.

3.7 Issue Statements

The sections above document cross-cutting themes that we identified throughout the entire Gen(Z)AI process, based on participant discussion. What follows outlines the specific issues that participants described as their most pressing concerns in each of the regional forums. These issues animate the policy recommendations outlined in section 4 of this report.

The issue statements below are the direct outputs of the deliberative process described in section 2, with specific language articulated and voted upon by participants. They describe issues that participants, having learned from experts and worked through structured discussion, collectively identified as the most consequential problems in each policy domain. Each was refined through facilitated small-group deliberation, tested against participants' own values and lived experience, and agreed upon by a minimum of 90 percent of forum participants before being carried forward.

Issue Statements from Forum #1 on AI Chatbots

1. Overreliance on AI chatbots can lead to emotional dependence that exacerbates social isolation, atomization, and contributes to mental health issues.
2. Widespread use of AI chatbots risks cognitive offloading and loss of critical thinking skills among users, jeopardising our abilities to learn, work, and engage in civic discourse.
3. AI chatbots increase users' potential exposure to harmful content, including sexually explicit, extremist, and/or self-harm content.

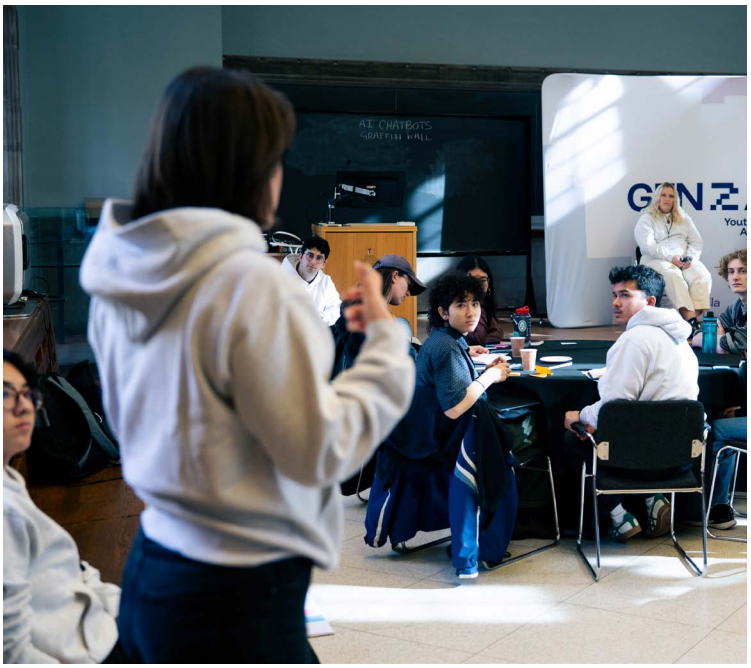

The Toronto Forum, photo by Tishan Baldeo

Issue Statements from Forum #2 on AI & Information Integrity

1. AI-generated content, including mis- and dis-information, overwhelms users and undermines confidence in reliable and accurate information, with disproportionate effects on vulnerable populations.
2. AI recommendation systems push ideologically extreme content and reinforce echo chambers, contributing to social and political polarization online and offline.
3. An information environment driven by engagement-based incentives and flooded with mis- and dis-information contributes to declining trust in news, government, and other institutions.

Issue Statements from Forum #3 on AI & Data Privacy

1. The way that AI systems collect user data and disclose that collection is deliberately opaque, making it difficult for users to provide informed consent about how their data is being used and sold. This may disproportionately affect vulnerable groups, including children and the elderly.
2. AI systems are not subject to adequate and enforceable safeguards for the collection, storage, and sharing of user data. This may lead to (un)intended harms, including, but not limited to, excessive surveillance and profiling.

Issue Statements from Forum #4 on AI & Age Assurance

1. AI-enabled age assurance technologies undermine privacy and digital safety by exposing users to potentially harmful storage and use of their data.
2. The use of verification technologies to restrict young people's access to social media, AI platforms, and online community spaces increases the risk that restricted users could be displaced to unmoderated and dangerous platforms.
3. Generative AI platforms pose unique risks to children, including by increasing their access to harmful content and negatively impacting their cognitive and social development.

4

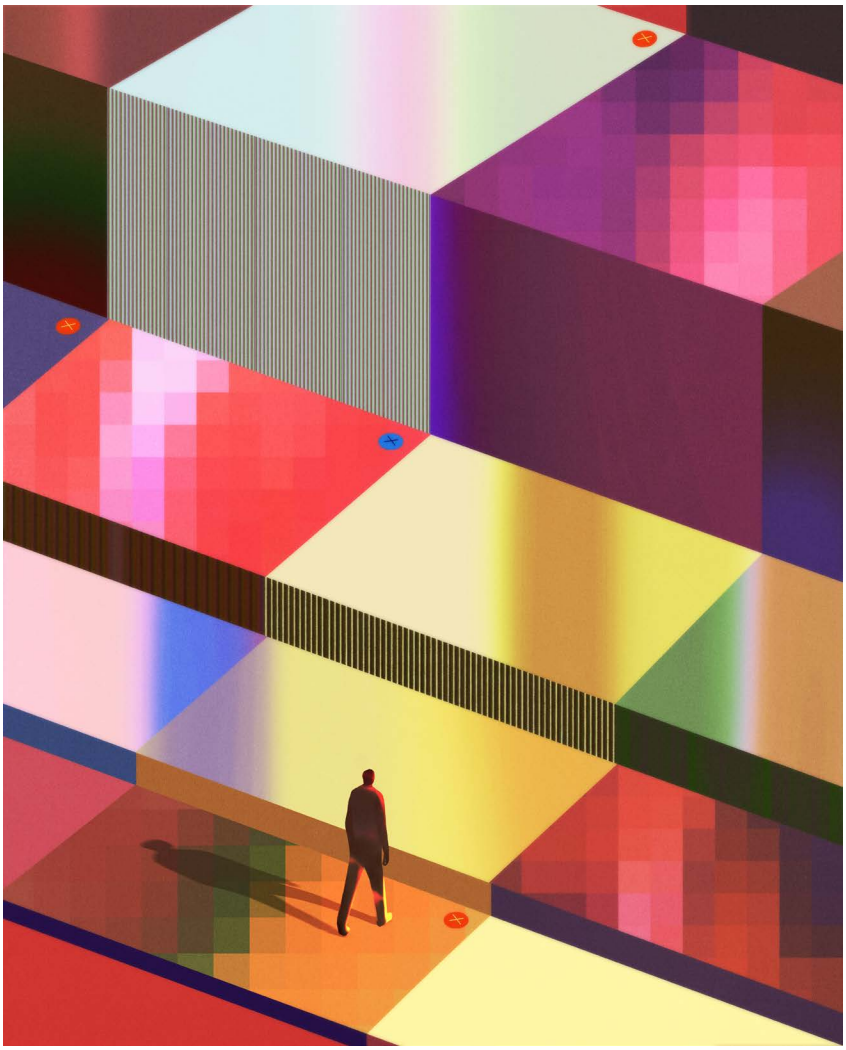

POLICY RECOMMENDATIONS

The following section presents the policy recommendations as written by Gen(Z)AI participants across the four regional forums.[7] These recommendations were produced through the structured deliberative process described in Section 2, and received a minimum consensus threshold of 90%. The issue statements that frame these recommendations are presented in the preceding section.

The recommendations span four policy domains at the intersection of AI governance and online harms: AI chatbots, information integrity, data privacy, and age assurance. Following the presentation of these recommendations, section 5 details relevant policy context from both Canadian and international jurisdictions and suggests how the recommendations written by the Gen(Z)AI participants might be operationalized federally in the context of the current political moment in Canada.

This context is drawn, in part, from the policy resources provided to participants in the lead up to each forum, and from the expert presentations that the participants received during each virtual learning day.

4.1. AI Chatbots

Gen(Z)AI participants developed five recommendations addressing the chatbot governance challenge:

1. Mandate that AI platforms address the addictive design of AI chatbots by requiring measures such as content filters and optional data cache deletion, and explicitly providing users with the ability to determine levels of responsiveness and conversationality.
2. Mandate accessible flagging capacity for users, require platforms to regularly report these instances, in a timely fashion, to an independent body with enforcement capacity, and make such reports accessible to the Canadian public.
3. Mandate that social media platforms and search engines have explicitly and easy-to-navigate opt-out options for integrated AI technologies, and compel platforms to provide users with explanations, in plain- and age-appropriate language, of the implications of opting in or out of AI integration.
4. Establish a funding program to support organizations offering critical AI literacy programming and resources to Canadians, including public consultations and deliberative engagements on AI's impacts that involve intergenerational and diverse groups of Canadians.
5. Establish a new, independent government body to enforce AI safety standards, conduct systems evaluations, algorithm audits, and risk assessments, and intake user complaints, including by offering dispute resolution and other resource mechanisms.

The Montreal Forum, photo by Jack Devine

4.2 AI & Information Integrity

Gen(Z)AI participants developed five recommendations addressing challenges at the intersection of AI and information integrity:

1. Mandate that digital platforms explicitly label AI-generated content and give users the functionality to omit this content.
2. Give people copyright over their own features and likeness, and create an online regulator to enforce the removal of non-consensual AI-generated material, including Child Sexual Abuse Material (CSAM).
3. Mandate that platforms monitor, flag, and transparently share information, with both the public and government, about the spread of mis- and dis-information, especially during high risk moments, including elections and public health crises.
4. Require that platforms introduce standards for AI-recommendation systems and data profiling processes to limit the spread of harmful content and known and suspected bot activity, and promote local content.
5. Mandate that platforms implement an independent third-party authentication mechanism for content posted from news and other public service organizations to promote and prioritize credible and reliable content.

4.3 AI & Data Privacy

Gen(Z)AI participants developed three recommendations addressing challenges related to AI and data privacy:

1. Mandate that platforms and AI companies implement meaningful and informed consent mechanisms to users, including by publishing a version of their terms and conditions that uses plain-language and is accessible by default.
2. Impose privacy-by-default standards for all AI systems.
3. Develop enforceable data privacy standards that require that platforms and AI companies:
    a. Prevent purpose-creep by strictly limiting data handling to clearly-defined and user consented purposes;
    b. Provide accessible options to delete user data upon request; and,
    c. Be subject to harsher legal consequences for the mishandling of users' sensitive data, including, but not limited to, health, financial and identity-related information.

4.4 AI & Age Assurance

Gen(Z)AI participants developed four recommendations related to AI and age assurance:

1. Create a standardized age-verification system to restrict users' access to generative AI platforms through the creation of an anonymized digital token system, with associated programs and accessible resources to inform the public about its implementation.
2. Mandate, in cases where age assurance is used, that companies adhere to stronger regulation, enforced by a Regulator, surrounding the use of sensitive age assurance data, including by:
    a. Imposing time-limited storage;
    b. Imposing safety audits on platforms and third-party data collectors; and,
    c. Requiring leakage protection in models and training.
3. Mandate that any AI platforms accessible to children, including in educational contexts, implement safety-by-design protocols to safeguard their use and promote learning and skills development.
4. Introduce measures to minimize young users' exposure to harmful design features and content across digital platforms, including by:
    a. Increasing transparency measures, by requiring digital safety and compliance reporting;
    b. Requiring that AI generated media content be identified through digital watermarking technology; and
    c. Reducing sycophancy and positive reinforcement in AI chatbots.

The Toronto Forum, photo by Tishan Baldeo

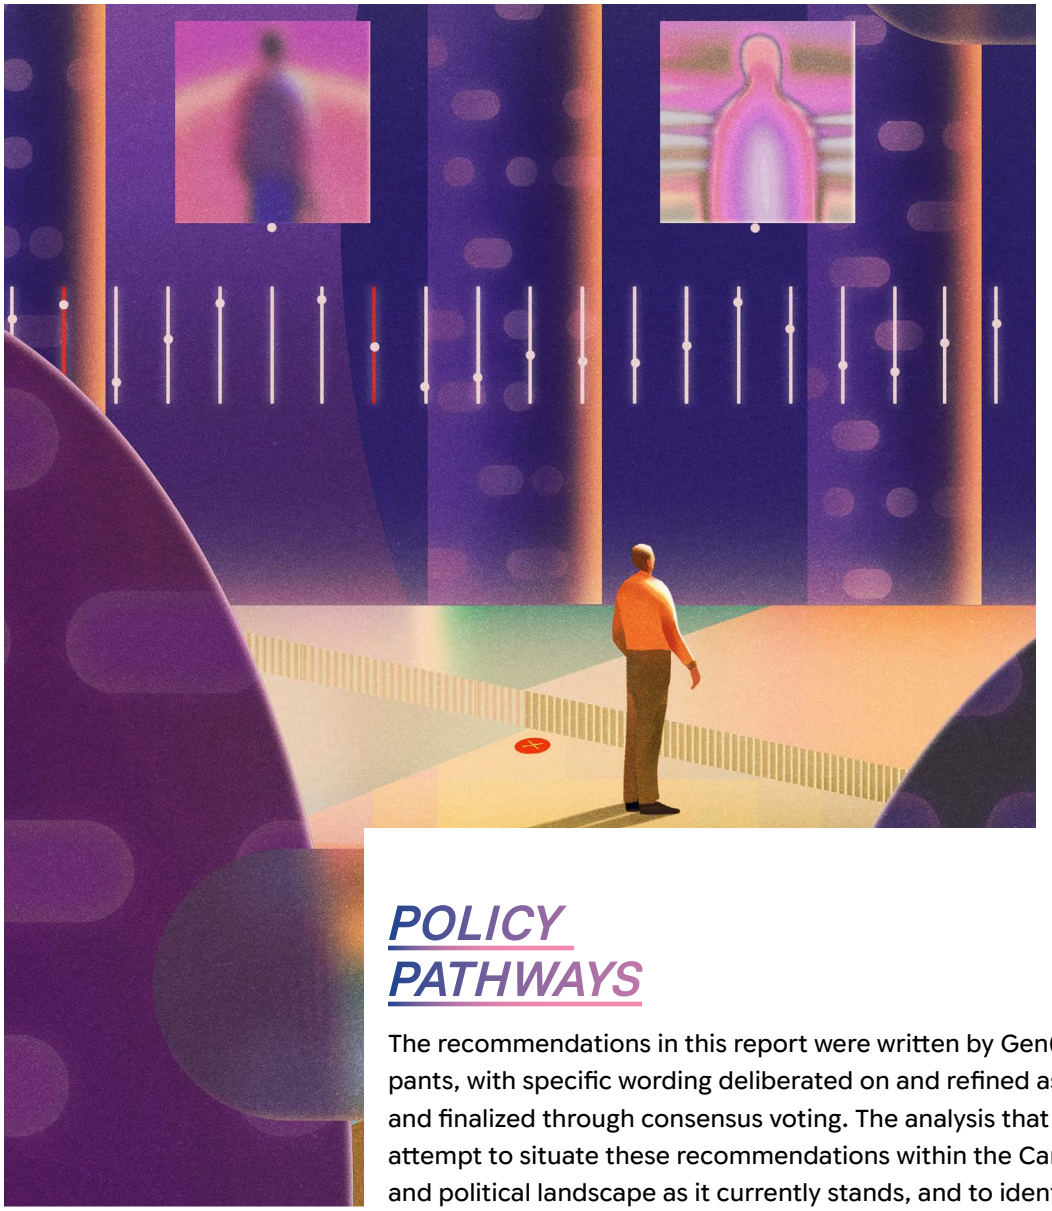

5

POLICY PATHWAYS

The recommendations in this report were written by Gen(Z)AI participants, with specific wording deliberated on and refined as a group, and finalized through consensus voting. The analysis that follows is an attempt to situate these recommendations within the Canadian policy and political landscape as it currently stands, and to identify a set of pathways through which they could be operationalized in the near term. We offer this as a demonstration that the recommendations that Gen(Z)AI participants drafted are concrete and actionable policy proposals that can support the development of two forthcoming pieces of legislation:

1. **Online Harms Legislation** that:
    a. Scopes in AI chatbots as a regulated technology;
    b. Centres safety-by-design and age-appropriate design; and,
    c. Props up an independent regulator with real enforcement capacity.
2. **Updated Privacy Legislation** that:
    a. Classifies children's data as sensitive by default; and,
    b. Extends protections to inferred and derived data.

This suggested policy pathway represents one shape that participants' recommendations could take when read against the current legislative landscape in Canada. We do not assume to account for every single proposal raised above, and encourage readers to consider this analysis as a complement to the recommendations that Gen(Z)AI's participants have crafted. This section offers a roadmap for near-term operationalization.

5.1 New Online Harms Legislation

A successor to Bill C-63 is the most direct legislative vehicle for addressing what Gen(Z)AI participants called for in their recommendations. When C-63 was tabled in February 2024, it placed substantive obligations on large social media services: 1) a duty to act responsibly; 2) a duty to protect children through age-appropriate design; 3) a duty to make certain content inaccessible; and, 4) a duty to keep records. It also would have propped up a dedicated enforcement architecture, including a Digital Safety Commission, a Digital Safety Ombudsperson, and a Digital Safety Office, with significant financial penalties for non-compliance.

The most important design decision for any new online harms legislation is how consumer-facing AI systems, including AI chatbots, are scoped in. C-63 defined a regulated "social media service" as a platform whose primary purpose is facilitating interpersonal communication between users, a definition that does not capture AI chatbots, whose primary relationship is between system and individual. To meet the concerns of Canadian youth as expressed through Gen(Z)AI, a C-63 successor should bring consumer-facing generative AI systems within its scope, defined by characteristics such as public accessibility, type of model interaction, and scale and reach, rather than by the interpersonal communication model that the Act originally used. If this were the case, the Digital Safety Commission's mandate would need to extend to cover these systems by virtue of the same scoping decision.

Participants called for age-appropriate design requirements to be specified in legislation, including content filters to address addictive design, user-controlled data cache deletion, adjustable levels of AI responsiveness, and digital watermarking of AI-generated content. These are structural policy responses that reflect a coherent analytical position: that the harms youth experience are, in their view, a product of deliberate design outcomes, and that legislation that addresses only content, and only after the fact, cannot reach the source of the problem. C-63's child protection duty established the right principle but deferred its content to regulation, a choice that prioritized regulatory flexibility over legislative specificity. Without design requirements specified in law, obligations are only as strong as the regulator's willingness and capacity to enforce them. A successor framework should weigh whether the nature of the harms expressed by Gen(Z)AI participants calls for the greater certainty that legislative specification would provide.

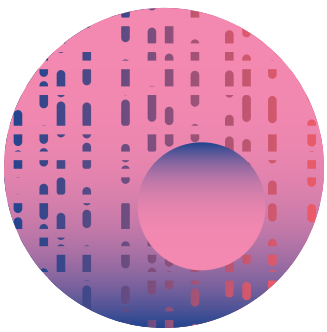

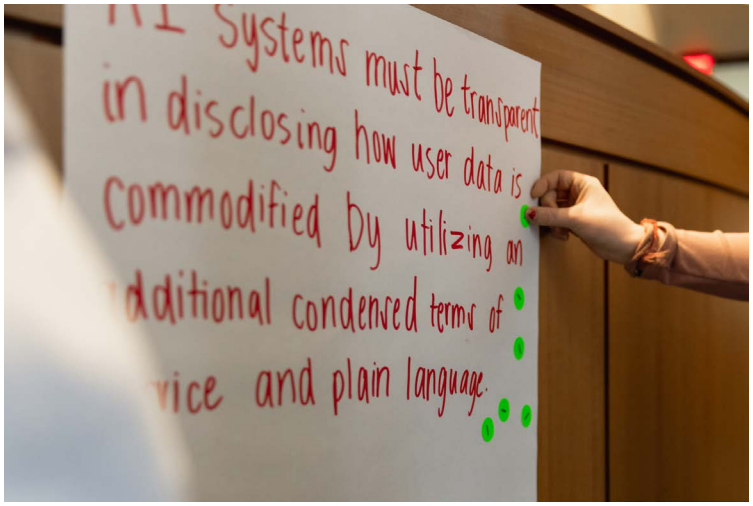

The Vancouver Forum, photos by Alison Boulier

To align with Gen(Z)AI's recommendations, the institutional architecture that C-63 proposed would likewise need to be extended in scope and capacity. C-63 gave the Commission investigative and inspection powers, including the ability to access electronic data inventories. Whether those powers are sufficient for the kind of proactive algorithmic oversight (i.e. standing risk assessments, ongoing audits of AI systems design, etc.) that participants called for is an open question, and one worth specifying in any successor legislation. Gen(Z)AI participants were explicit that a regulator would need genuine technical expertise to assess age assurance systems, evaluate third-party data collectors, and manage the implementation of a digital token system, should age verification be implemented as an online harms measure. That is a different institutional profile than a complaints-based enforcement body, and the mandate and resourcing should reflect it.

Preserving C-63's Ombudsperson function would support Gen(Z)AI's identification of recourse as a structural problem. An Ombudsperson with a genuine public interest mandate gives avenues for redress that are currently limited and largely controlled by platforms themselves. The alternative—assigning these functions to the CRTC, whose primary mandate is broadcasting and telecommunications—would risk doing precisely what participants argued against: bolting AI and online harms oversight onto an institution whose expertise and focus lie elsewhere. Extending the institutional design and scope of C-63 seems to be the most coherent model to move forward with many of Gen(Z)AI's policy recommendations.

5.2 An Updated Privacy Framework

Canada's existing privacy framework is not adequate to the environment that Gen(Z)AI participants described. PIPEDA was designed for a pre-AI commercial world and contains no classification of data by sensitivity, no provisions for inferred or derived data, and a consent model built for transactional data relationships rather than continuous ones. Bill C-27's Consumer Privacy Protection Act (CPPA) was the most serious legislative attempt yet to close these gaps. Several of its provisions moved directly in the direction that Gen(Z)AI's participants asked for. For example, it classified the personal information of minors as sensitive by default, a provision that is both foundational and, if carried forward in a successor framework, would provide the legal basis for the stronger child protections that Gen(Z)AI participants proposed. The CPPA also introduced enhanced consent provisions, including prohibitions on consent obtained by deception and a right to withdraw consent; purpose limitation requirements that confined data handling to appropriate and specified uses; a right to disposal of personal information at an individual's request; and plain-language access requirements, which would partially address what Vancouver participants identified as a core accessibility failure of existing consent mechanisms.

Participants asked that children's data be treated as sensitive when collected and that AI systems interacting with children be designed from the outset to minimize collection, restrict use to clearly-consented purposes, and protect children from the specific harms that flow from how those systems are built. The Vancouver Forum's recommendations for privacy-by-default standards and enforceable purpose limitation, and the Halifax Forum's recommendations for time-limited storage of age assurance data, mandatory safety audits of third-party data collectors, and protections against leakage in AI model training, all require design obligations that the CPPA's consent-based framework does not establish, but could in a renewed version. Privacy-by-default has clear international precedent in the EU's General Data Protection Regulation, which mandates it as a condition of lawful processing.

The most significant policy gap concerns inferred and derived data. The CPPA governs personal information that organizations collect, use, or disclose, but it does not explicitly address the data that AI systems generate about individuals from their behaviour. To meet Gen(Z)AI participants' recommendations, a CPPA successor should extend its protections explicitly to inferred and derived data, treating outputs generated about an individual as personal information subject to the same consent, access, purpose limitation, and disposal rights as directly collected data.

The Vancouver Forum, photos by Alison Boulier

Conclusion

FROM GOVERNANCE FOR YOUTH TO GOVERNANCE WITH YOUTH

Throughout Gen(Z)AI, one hundred young Canadians, given access to adequate information and structured deliberative space, produced coherent, thoughtful and actionable policy recommendations. The project's "layered deliberation" model showed that this deep, in-person deliberation and large-scale digital engagement are not mutually exclusive, but can be combined to produce democratically legitimate input at scale. Participant reflections (Appendix 1) further underscore the capacity of processes like this to not only produce evidence-based outputs, but also transform participants' sense of civic agency and political possibility.

More broadly, Gen(Z)AI demonstrates the value of co-governance. By mobilizing the experiential knowledge of those most affected, processes like this can produce more grounded and responsive policy, while also building public trust in both government and the possibility of a safer, more accountable digital ecosystem. This is not a trade-off between innovation and protection, but a foundation for both.

From coast to coast, the dedicated and nuanced work of Gen(Z)AI's participants has demonstrated that young people are ready to lead on this file. The question is whether the rest of us, and Canadian institutions more broadly, are ready to listen.

APPENDIX

Appendix 1: Participant Testimonials

"Gen(Z)AI allowed me to engage deeply with relevant issues of my generation. I am proud of the work we have achieved as a group, and hope that policymakers will listen to our collective voice with intent."

"This was such an insightful project to be a part of! It was great to feel like we, as youth, had a say in developing policies that would directly affect us. Oftentimes, and especially with everything that's going on in the world, it's hard to feel like anything that we say or do is actually being considered/has any impact. I also appreciated how the participants came from a variety of academic and life backgrounds to better inform the discussion, since we all had varying perspectives."

"I really had a great time going through the process and meeting so many young people my age engaged in the topic!"

"I feel that youth, in addition to all other Canadians, need a chance to engage in processes like Gen(Z)AI to ensure we understand the technologies that will impact our futures and have the opportunity to share our hopes, fears and desires for how these technologies will be integrated into our nation's fabric."

"It is essential to gather the input of a broader share of the population given that decisions about AI development and adoption are being made by a limited number of people; which is arguably undemocratic, as these technologies will radically reshape all of our lives."

"Gen(Z)AI was a good forum for giving young people a voice on AI. I believe that, just as with other issues like the environment, health, and education, the voices of young people must be heard. I even believe that when it comes to artificial intelligence, young people need to make themselves heard as much as possible. The fact that AI has gone from simple algorithms used to run computers or non-player characters in video games to LLMs, algorithmic models predicting financial markets, systems that are close to understanding or simulating consciousness, and increasingly advanced robotics, is for me both fascinating and terrifying. Yet this experience makes me feel that deep down, we are all in this together, with similar concerns. We simply need to work together: trust the experts conducting research on the implementation of ethical AI, listen to civil society, put regulations in place, collaborate with tech companies and their shareholders, and together, we will get there."

"Gen(Z)AI was an extraordinary opportunity not only to engage as a youth but to deliberate, discuss, and debate issues that we feel passionate about. Too often, young voices feel unheard or dismissed; however, this process gave us not only the tools but also the ability to share our opinions and perspectives in a meaningful way, with not only our peers but policymakers as well. It is an important reminder that democracy is a shared value and a fundamental right for people of all ages, and I'm glad to have been a part of it."

Appendix 2: List of Participants

Gen(Z)AI was made up of the following participants, from whom we are grateful to have received such dedicated and thoughtful participation:

Toronto Cohort: Alexandra Farrugia, Alexandra Wilson-Zegil, Alexandre Krantz, Alonso Munoz Sanchez, Danylo Ustymenko, Emanuel Sankar-Vanderherberg, Emily Chen, Faridah Amoo, Georgia Berg, Graeme Lavrence, Indigo Kim, Lakshna Ponrajah, Laryssa McGlashon, Majda Lojpur, Max Zuercher, Maxine Bisera, Morgan Bulitz, Raneem Zaitoun, Rubaina Farin, Santiago Carvallo Herrera, Sierra Wyllie, Simon Monterroza Prieto, Yusuf Zawahir

Montreal Cohort: Adanna Okorie, Adele Parts, Ameline Ramesan, Amine El Idrissi, Annie Carignan, Assia Harchaoui, Carla Zuercher, David O'Sullivan, Denali Tran-Le, Emily Hong, Étienne Langlois, Fabrice Blais-Savoie, Gracie Diabo, Jovan Rohac, Leo Litke, Mathis Larouche, Matthieu Martin du Pont, Nicole Trif, Nina Popovic, Noah Gregory, Patrick Brennan, Quinnlan Guest, Sophia Micomonaco, Sophia Rocco, Teerka Baskaran

Vancouver Cohort: Calvin Slinn, Cole Swetlishoff, Daniel Anene, Daniel Sass, Daphne Tian, Elling Speirs, Emily Hernandez, Ethan Fukuhara, Evan Minick, Ezra Scholl, Holden Sullivan, Joie Kathryn Marin, Leo Wu, Liam McKay-Argyriou, Lily Woods, Lucca Parolo, Maria Kapitanenko, Nicole Brusilovskaya, Partap Grewal, Sloane Madden, Solenne Le Billon, Steewa Maria Philip, Zoe Stojanovic

Halifax Cohort: Abby Attree-Cannon, Adeb Arianson, Allie Luscombe, An-Kita Das, Austin O'Neill, Cameron Tynes, Chinazam Igwe, Eniola Oyegunwa, Gift Iluyomade, Hayden Sparkes, Jack Lewandowski, Jamal Raaki, Jana Salem, Jinlin Tan, Karthikeya Redrowtu, Kate Dobbins, Kayleigh Stuart, Laura Bremner, Lister de Vitré, Mackenzie Tapp, Mersn Maniania, Peter Oyedijo, Rebecca Lawrence, Samantha Sherman, Sophie Baker, Vatsal Ichplani, Wells Mayes

Appendix 3: List of Expert Speakers

Toronto Forum #1

- **Kris Perry,** Executive Director, Children and Screens
- **Meetali Jain**, Founding Executive Director, Tech Justice Law
- **Taylor Lynn Curtis**, Software Developer, Mila
- **Maddie Freeman**, Founder and Executive Director, NoSo Connection Collective

Montreal Forum #2

- **Alex MacIsaac**, Senior Research Coordinator, Samara Centre for Democracy
- **Melanie Walker**, Executive Director, Canadian International Council
- **Halyna Padalko**, Fulbright Scholar, Massachusetts Institute of Technology
- **Aengus Bridgman**, Associate Director (Research), Centre for Media, Technology, and Democracy

Vancouver Forum #3

- **Teresa Scassa**, Canada Research Chair in Information Law and Policy, University of Ottawa Law
- **Kris Klein**, President, IAPP Canada
- **Sara Grimes**, Wolfe Chair in Scientific and Technological Literacy, McGill University
- **Riley McNair**, Senior Advisor, Information and Privacy Commissioner of Ontario
- **oline Twiss**, Deputy Information and Privacy Commissioner, Office of the Information and Privacy Commissioner for BC

Halifax Forum #4

- **Emily Laidlaw**, Canada Research Chair in Cybersecurity Law, University of Calgary
- **Michael Geist**, Canada Research Chair in Internet and e-Commerce Law, University of Ottawa
- **Marie-Eve Nadeau**, Head of International Affairs, 5Rights
- **Alex MacIsaac**, Senior Research Coordinator, Samara Centre for Democracy
- **Michael Karanicolas**, James S. Palmer Chair in Public Policy & Law, Dalhousie University

Appendix 4: Gen(Z)AI Youth Fellows and Project Leadership Bios

Helen A. Hayes

Helen A. Hayes is the co-lead of Gen(Z)AI. Alongside Fergus Linley-Mota, Helen designed the project's methodology, trained its facilitators, and acted as the lead author of the project's final report. Hayes is a Fellow at both Mila and the Morris J. Wosk Centre for Dialogue, and is a 2026 recipient of Forbes 30 Under 30 (Education) for her work engaging young people in digital policymaking, as well as for her teaching in undergraduate and graduate programs at McGill University. Her work has been published in the Harvard Journal of Law and Technology, Interfaces, Election Law Journal, and Lecture Notes in Computer Science, among others. She is the co-author of Voting Online: Technology and Democracy in Municipal Elections (McGill-Queen's University Press, 2024), and a co-editor of Regulating Digital (University of Toronto Press, 2026) and Trust and Trustworthiness: Evaluations of Online Voting (McGill-Queen's University Press, 2026).

Fergus Linley-Mota

Fergus Linley-Mota is the co-lead of Gen(Z)AI. Alongside Helen A. Hayes, Fergus designed the project's methodology, trained its facilitators, and acted as the lead author of the project's final report. Fergus is the Founder and Director of the Dialogue on Technology Project, housed at Simon Fraser University, where his work focuses on expanding the range of voices shaping technology governance. He is also a Senior Fellow at the Centre for Media, Technology and Democracy at McGill University. He has degrees from SFU and the University of Oxford, and is a 2025/26 Action Canada Fellow.

Madeleine Case

Madeleine Case is the Lead, Policy and Strategic Initiatives at the Democratic Engagement Exchange at Toronto Metropolitan University. Her work focuses on strengthening democratic participation through partnerships, cross-sector collaboration, and research. A former federal political staffer, Maddie brings experience from government to her work on democratic engagement. Maddie is also a Youth Fellow with the Centre for Media, Technology and Democracy at McGill University, focusing on youth engagement and the evolving digital information environment. She holds a Master's of Public Policy and Administration from Toronto Metropolitan University and a Bachelor of Arts in Political Studies from Queen's University.

Alexander Martin

Alexander Martin is a Youth Fellow at the Centre for Media, Technology and Democracy and a Digital Policy Hub doctoral fellow at the Centre for International Governance Innovation (CIGI). He is a Ph.D. candidate in science and technology studies at York University, where his research focuses on information integrity, digital governance, and digital communication infrastructure in Canada. He is a deputy editor at Open Canada, part of the Canadian International Council (CIC), where he works with authors on Canadian foreign policy and technology policy issues. He previously held policy and strategic communications roles at the North Atlantic Treaty Organization's (NATO) Office of the Chief Scientist and the United Nations Information Integrity Unit, focusing on mis- and disinformation and hate speech in peace operations. His work examines how public-interest digital infrastructure and governance can strengthen democratic resilience and digital sovereignty in Canada.

Julian Lam

Julian Lam is a JD candidate at the Peter A. Allard School of Law at UBC. His work focuses around the confluence of law, technology, and public policy, with a focus on AI governance, venture, and international trade. He is a Youth Fellow with the Centre for Media, Technology and Democracy (MTD), where he has contributed to national and global initiatives on AI governance and also a digital policy hub fellow at the Centre for International Governance Innovation (CIGI). Julian brings experience across public, private, and innovation sectors, including roles at CPP Investments, the Vancouver Fraser Port Authority, Arc'teryx, and Canada's Digital Technology Supercluster, where he helped lead the development of the latter's global strategy. His current work explores the governance of emerging technologies and the role of law in shaping equitable, globally connected technological futures.

Nonso Morah

Nonso Morah is a McCall MacBain Scholar and graduate student in the Faculty of Education at McGill University. Her research explores community-based knowledge production, digital youth empowerment, and diasporic education. She holds an Honours Bachelor of Social Sciences from the University of Ottawa and has previously worked as a federal political staffer and non-profit coordinator with the 1834 Foundation. Having held several positions across the public, nonprofit, and arts sectors, Nonso brings an expansive, interdisciplinary perspective to fostering youth civic engagement across Canada. As a Youth Fellow with the Centre for Media, Technology and Democracy, she contributed to the Gen(Z)AI initiative and supported the Centre's digital media strategy. Beyond academia, Nonso is a published poet and award-winning spoken word artist who has performed for the National Gallery of Canada, RogersTV, and the Government of Canada.

Appendix 5: Advisory Board Members and Bios

Blair Attard-Frost

Blair Attard-Frost is an Assistant Professor in the Department of Political Science and a Fellow at the Alberta Machine Intelligence Institute (Amii). She completed her PhD, Master of Information (MI), and Honours Bachelor of Arts (HBA) at University of Toronto. Blair's research applies a trans feminist lens to address challenges of power, participation, and justice in the governance of artificial intelligence. Her research appears in academic journals such as *Big Data & Society*, *First Monday*, *Government Information Quarterly*, *AI and Ethics*, and *International Journal of Information Management*.

Alicia Combaz

Alicia Combaz is the Founder and CEO of Make.org, a pioneering civic tech organization in Europe. Before launching Make.org, she spent a decade in the tech industry, where she built and led teams of engineers, data scientists, and product managers. Notably, she established France's first Growth Hacking Team at Deezer, driving the company's rapid expansion. In 2016, Alicia co-founded Make.org with Deezer's former CEO, aiming to use technology to strengthen democracy. Her work has been recognized through features like Le Point-Giverny's 2021 list of outstanding individuals under 40 and her selection as a Rising Talent by the Women's Forum for the Economy & Society in 2022. In December 2025, she was awarded the prestigious title of Chevalier de l'Ordre National du Mérite by the French government.

Aftab Erfan

Dr. Aftab Erfan is the Executive Director of the SFU Centre for Dialogue and Associate Member at SFU School of Public Policy. Previously, she served as the City of Vancouver's inaugural Chief Equity Officer, and Director of Dialogue and Conflict Engagement at UBC. She holds a PhD in planning from UBC, a Masters in planning from McGill University, and a BSc in environmental sciences from UBC. Aftab has extensive experience as a process designer and facilitator, specializing in hosting dialogues on contentious issues with warmth and creativity. Her experience spans 15 years and five continents. She holds numerous honours, including BC500 most influential business leaders in BC (Business in Vancouver), Research for a Better Life: The Storytellers Challenge (Social Sciences and Humanities Research Council), and Best Published Paper Award Shortlist (Association of European Schools of Planning).

Sara Grimes

Sara M. Grimes is the Wolfe Chair in Scientific and Technological Literacy, founding Director of the Kids Play Tech Lab, and a Full Professor in the Communication Studies at McGill University. Her research centers on understanding emerging trends and issues related to the technologization of childhood and supporting the development of ethical designs, business practices, and policies that advance children's rights and wellbeing in the digital environment. Her current and upcoming projects include examinations of age-appropriate design in games, the legal and ethical dimensions of algorithmic marketing to children, as well as a multiyear study of children's creativity, play and parasocial relationships with AI in digital games. She is the author of the award-winning book, *Digital Playgrounds: The Hidden Politics of Children's Online Play Spaces, Virtual Worlds, and Connected Games*, and the upcoming *Kidfluenced: Children's Digital Media-Making, Content Creation, and Cultural Rights*.

Sally Guy

Sally Guy is the Director of Public Policy and Inclusion at Mila, where she leads projects and programs at the intersection of AI, public policy, and inclusion—all aimed at advancing Mila's mission to ensure that AI innovation serves the benefit of all. Previously, Sally served as Director of Policy and Strategy at the Canadian Association of Social Workers (CASW). There, she shaped policy aimed at addressing socioeconomic inequities and built relationships with governments and civil society actors to advance a more just and equitable Canada. Her other past affiliations include sitting on the Board of the Canadian Alliance on Mental Illness and Mental Health and the National Guidance and Implementation Committee for the VEGA Project (a 10-year, $100M federal initiative supporting victims of family and intimate partner violence).

Ness Kenalty

Ness Kenalty is the Director of Childhood Development at Waltons Trust, a Canadian grant-making foundation, where she leads a granting portfolio focused on reducing excessive and harmful screen time among children. She holds a degree from the University of Toronto's Ontario Institute for Studies in Education and a Master's in Early Childhood Development from Toronto Metropolitan University. Prior to joining Waltons, Ness worked in both schools and nonprofit organizations, giving her a well-rounded perspective on child development, education, and policy.

Marie-Ève Nadeau

Marie-Ève Nadeau is Head of International Affairs at the 5Rights Foundation, leading global efforts to embed children's rights, safety and privacy by design and default in technology governance. Based in Brussels, she has spent the last five years collaborating with governments, regulators, international institutions and civil society organisations to strengthen AI, data protection and online safety policy frameworks worldwide. She has a background in international law and human rights. Before joining 5Rights, she promoted human rights internationally through her work with the EU office of the International Federation for Human Rights (FIDH), the United Nations Development Programme (UNDP), and Québec Without Borders in Peru.

Sam Reusch

Sam Reusch is the Executive Director of Apathy is Boring, a leading non-partisan democracy organization mobilizing young people across Canada. Since 2020, she has led national growth, scaled year-round programs, and delivered record-setting youth participation and turnout campaigns. Her work centers on democratic reform, elections, and countering mis- and disinformation, with a focus on rebuilding trust and strengthening participation. Sam is a recipient of the King Charles III Coronation Medal and serves on the board of Volunteer Canada. Outside of work, Sam enjoys good food, long walks, and spending time with books and podcasts.

Valerie Steeves

Valerie Steeves is a Faculty Member at the Centre for Law, Technology and Society and a Full Professor in the Department of Criminology of the Faculty of Social Sciences at the University of Ottawa. Steeves' main area of research focuses on the impact of new technologies on human rights. She is the principal investigator of The eQuality Project, an SSHRC-funded partnership of researchers, educators, advocates, civil society groups, and policymakers interested in examining the impact of online commercial profiling on children's identities and social relationships. As the lead researcher for MediaSmart's Young Canadians in a Wired World research project, she has tracked young people's use of new media from 2004 to 2020.

Endnotes

1 The Gen(Z)AI Advisory Board brought together world-leading scholars, practitioners, and advocates: Blair Attard-Frost, Alicia Combaz, Aftab Erfan, Sara Grimes, Sally Guy, Ness Kenalty, Marie-Eve Nadeau, Samantha Reusch, and Valerie Steeves. Their bios are included in Appendix 5.

2 A civic lottery is a balanced way of selecting the members of a citizens' assembly. It is based on a form of sortition that uses Canadian census data and a randomized selection process to recruit participants. The result is a group of volunteers that broadly match the demographics of the jurisdiction they represent. This stratified sampling methodology ensures that members of Gen(Z)AI were selected at random, but in a way that broadly represented the demographics of the Canadian provinces in which forums were held, balancing for gender, sexual orientation, ethnic identity, and Indigeneity.

3 The expert speakers for each forum are listed in Appendix 3.

4 The Forums were held at Victoria College in the University of Toronto (Toronto Forum #1), The McCall MacBain Faculty of Arts at McGill University (Montreal Forum #2), The Morris J. Wosk Centre for Dialogue at Simon Fraser University (Vancouver Forum #3), and the MacEachen Institute at Dalhousie University (Halifax Forum #4).

5 The Gen(Z)AI youth fellows are: Maddie Case, Alexander Martin, Julian Lam, and Nonso Morah. Their biographies are available in Appendix 4.

6 The Gen(Z)AI project leads are Helen A. Hayes and Fergus Linley-Mota. Their biographies are available in Appendix 4.

7 These policy recommendations were informed by a "values-mapping" exercise, wherein participants were asked to surface the values that they hoped to guide their policy development phase. These included: Curiosity; Authenticity; Balance; Connection; Trust; Community; Wellbeing; Responsibility; Making a Difference; Reliability; Security; Future Generations; Accountability; Optimism; Fairness; Stewardship; Integrity; Adaptability; Safety; Service.