RESPONSIBLE AI AND CHILDREN: INSIGHTS, IMPLICATIONS, AND BEST PRACTICES

SARA M. GRIMES
ALISSA N. ANTLE
VALERIE STEEVES
NATALIE COULTER

APRIL 2024

CIFAR

AI Insights

ABOUT THE AUTHORS

SARA M. GRIMES, PHD

Professor, Faculty of Information
Bell University Labs Chair in Human-Computer Interaction
University of Toronto

ALISSA N. ANTLE, PHD

Professor, School of Interactive Arts & Technology
Royal Society of Canada College of New Scholars
Simon Fraser University

VALERIE STEEVES, PHD

Professor, Department of Criminology and the Centre for Law, Technology and Society
Co-Lead, The eQuality Project
University of Ottawa

NATALIE COULTER, PHD

Associate Professor, Department of Communication & Media Studies
Director, Institute for Research on Digital Literacies
York University

LAND ACKNOWLEDGMENT

We wish to acknowledge this land on which CIFAR operates. For thousands of years it has been the traditional territory of many nations including the Mississaugas of the Credit, the Anishnabeg, the Chippewa, the Haudenosaunee and the Wendat peoples and is now home to many diverse First Nations, Inuit and Métis peoples. We are grateful to have the opportunity to work on this land.
We also acknowledge we are all responsible for reconciliation. CIFAR's AI & Society program seeks to advance our understanding of the societal implications of AI to design a future of responsible AI. A future of responsible AI includes one that centres the concerns of Indigenous communities. CIFAR is committed to prioritizing Indigenous perspectives in the development and design of responsible AI.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

In this policy brief, we distill findings from relevant previous research that provides much needed insight into why and how children use data-centric technologies in their everyday lives, play and learning, and how this impacts their social relationships, development and well-being, privacy and other rights.

We comment on the various challenges and potential benefits of artificial intelligence (AI) for children, in light of prominent trends shaping children's access to, experience of, and relationship with emerging and antecedent data-driven technologies, and existing industries and systems.

We describe key insights produced from a systematic review of AI design research applying child-centric methodologies, an approach that is widely promoted in the literature but rarely explained in practical terms. We review relevant existing Canadian policy and provide recommendations for a rights-based approach to regulation moving forward.

We establish a timely child-centric and rights-based framework for thinking about responsible AI and children, that is applicable across technological devices and innovations and adaptable to diverse contexts of childhood.

We conclude with a set of key takeaways meant to guide future dialogue, design, and policy development in this crucial area.

1.0 INTRODUCTION

The increased integration of AI in critical areas of children's lives, from schools and social services to homes and social media platforms, means that children are already interacting with these technologies in multiple and complex ways (Ito et al., 2023). As AI technologies evolve and spread, they are more deeply integrated into children's lives–often in ways that are hidden or don't appear as AI, as in the case of "For You" feeds on TikTok (Grandinetti, 2023). The need for research on these impacts is clear, as is the need for proactive regulation and ethical industry standards that support the opportunities AI provides children while mitigating its risks.

A key challenge in writing this policy brief is the inconsistent way the terms AI, data-centric technologies, social media, and digital technologies are used in and across disciplines. For instance, the term AI is frequently used to describe applications that are data-centric or driven by predictive algorithms, but not "intelligent" (McEwen, 2023). Other times, AI refers to an anticipated future iteration. Instead of seeing this as a barrier, we approach it as an opportunity to make vital connections between interrelated technological forms, from long-established online social networks to newly-launched generative AI (GenAI) tools. In this brief, we use the term "data-centric technologies" to describe the growing array of applications, systems, and devices that collect, publish, process, analyse, and mobilize user data that now dominate our information society. This includes AI in its many iterations, but also associated and antecedent technologies, systems and processes.

Proposed new regulations, applications of existing regulations, and industry guidelines for ensuring "responsible AI" are emerging at rapid speed, in Canada and around the world. To date, however, children have been largely omitted from these broader policy discussions—or mentioned only briefly as users vulnerable to harm[1]. For example, while the current draft Bill C-27 includes special protections for the collection, use, and disclosure of personal information of minors, its section on AI (the Artificial Intelligence and Data Act (AIDA)) only refers to children once, in the Companion document, as an example of a "more vulnerable group."

Globally, 1 in 3 internet users are under the age of 18 years. In 2021, the United Nations confirmed that the Convention on the Rights of the Child (UNCRC) applies to the digital environment through its adoption of General Comment 25 (GC25). Yet, questions about how to keep AI "safe" for children are often met with blunt responses like age restrictions. As AI is integrated into more areas of children's lives—through increasing interaction with AI-driven applications, devices, and spaces—excluding child users becomes increasingly unethical. The need to address children's rights and best interests vis-à-vis AI is thus critical.

In this brief we take the position that initiatives and policies aimed at regulating and developing responsible AI must:

1. Consider the presence of children from the outset, while addressing their rights and best interests;

2. Ground any decisions or recommendations in both emerging research and the substantial existing literature on children's uses and relationships with antecedent data-centric technologies; and

3. Include children and adolescents in the research and development of AI technologies.

[1] An important exception is the European Union's EU Artificial Intelligence Act, which forbids the use of AI to manipulate children and confirms that an AI system's potential to negatively impact children's rights will be a factor in determining its risk classification.

2.0

CHILDREN AND DATA-CENTRIC TECHNOLOGIES

To date, the research on children and AI has largely focused on educational applications and children's attitudes towards AI (e.g., Hiniker et al., 2021). We still know very little about the extent to which children's behavioural data, personal information, and creative works are used to train AI, or about how diverse children are engaging (or not) with AI "in the field," outside of research contexts. The lack of ethical standards for using children's data in AI research and development, combined with longstanding disparities in how children and age are represented in the digital environment, introduces a risk of "age-related algorithmic bias" (Muralidharan et al, 2023). There is a clear need for more research in this area.

At the same time, the newness of the current generation of AI systems and tools is often overestimated. This can result in an ahistorical and decontextualized picture. While the recent breakthroughs in AI are significant, developing standards and policies to ensure that AI is made and managed responsibly requires an understanding of the techno-social contexts from which AI applications, companies, and uses emerge.

There is a substantial amount of relevant existing literature to draw on for understanding these contexts and histories. This includes a font of research on children's experiences with antecedent and related data-centric technologies. For over two decades, scholars from various fields have investigated the social, ethical, and developmental impacts of different data-centric technologies on diverse groups of children and adolescents. This literature maps the historical, social, and political-economic conditions out of which current and future iterations of AI are born.

The sections that follow provide a critical synthesis of relevant literature in four priority areas. It includes works from both specialized fields (e.g., child-computer interaction, and children's media studies) and traditional disciplines (e.g., sociology, and legal studies). It tracks children's use of data-centric technologies before and after the integration of AI. It provides a launchpad for improved dialogue between emerging and existing research, policymaking, and technology development going forward.

There is a tendency to overgeneralize when talking about "children," as policymakers and researchers extrapolate findings about one age group (often adolescents or even young adults) to "all" children. The term "child" is also often used in ways that affect race, class, and gender, even though historically much of the research has focused on middle class white boys. Drawing on the work of Konstatoni and Emejulu (2017), we acknowledge age as a vector of intersectional identity. Throughout this report, we refer to children as people aged 6 to 12 years, younger children (early childhood) as people aged 0 to 5 years, and adolescents as people aged 13 to 19 years. When discussing children's rights, we follow the definition used in the UNCRC of everyone under the age of 18 years.

2.2

EVERYDAY LIFE, PLAY, AND LEARNING

Adolescents and children are often among the earliest adopters and heaviest users of data-centric platforms, apps, and devices (Ito et al., 2010). A majority (86%) of Canadian children aged 9 to 12 years currently have at least one account on a platform such as TikTok or Snapchat (MediaSmarts, 2022). Messenger apps are used to communicate with peers; social media is used to make and share content; and games are used to hang out and have fun with friends and family. Young children play with apps that collect massive amounts of data from them, stream videos on YouTube, and connect with grandparents over Zoom. Children of all ages encounter data-centric technologies across multiple areas of their everyday lives (Mascheroni & Siibak, 2021).

Most Canadians have Internet access, but digital divides continue to disadvantage children in rural, racialized, and Indigenous communities due to disparities in connection quality, skills and literacies, and device capabilities (Helsper, 2021). Racialized and otherwise marginalized children are routinely subjected to biased algorithmic profiling, discrimination, and other harms through these systems (Noble, 2018). The broad integration of AI by social institutions, governments, and technological infrastructures means that children will be impacted by these systems in meaningful ways regardless of individual access or use patterns.

For many children and adolescents, data-centric technologies are important sources of information. In one study, US adolescents reported primarily using social media when seeking health information (Stevens et al., 2017). Increased access to information is beneficial and essential, but also carries risks of exposure to misinformation, disinformation, and extremist propaganda (e.g., Costello et al., 2020).

Many of the platforms and devices that children, young children, and adolescents use or come into contact with in their everyday lives contain AI at some level. Examples range from YouTube's recommender system and Roblox's content creation Assistant to smart security systems and fitness trackers (Pangrazio & Mavoa, 2023; Antle & Kitson, 2021). Children also engage with AI directly. One survey found that 91% of US households with children aged 2 to 8 years used conversational agents such as Apple's Siri or Amazon's Alexa, in over half (59%) of which children interacted directly with the conversational agent (Wronski, 2019).

2.2.1 PLAY

Most young children, children, and adolescents play digital games (Grimes, 2021). Some play with smart toys, such as social robots (Mascheroni & Holloway, 2019). Digital games and smart toys are subjects of controversy in the news and in policy debates, but the research is less divided. There is no scientific evidence that digital games cause "real-world" violent behaviour, or that children prefer playing with AI more than with other humans (Aguiar, 2021). Instead, the literature suggests that under certain conditions, digital play can benefit children's learning, identity formation, and well-being (Kafai & Fields, 2014; Grimes, 2021; Giddings, 2014).

Some digital games contain "persuasive designs," nudge techniques, or "dark patterns"—elements that draw on behavioural science, user data, and predictive algorithms to manipulate users into doing things they don't want to do. For example, Radesky et al. (2022) found that 65% of mobile game apps played by young children (aged 3 to 5 years) in their study contained features specifically designed to prolong gameplay, such as pop-up messages "from" the game's characters that appeared when children tried to quit the game, pressuring them to keep playing.

Persuasive design tactics are deployed for various reasons—from selling products to radicalizing players—and undermine many of the benefits of digital play, especially those associated with "free play" (Livingstone & Pothong, 2022). The potential that AI will be deployed in ways that amplify dark patterns is a growing source of concern (Mascheroni & Siibak, 2021; Willis, 2020). For example, "emotional AI"-driven toys that use biometric and behavioural data to assess and manipulate children's emotions are vulnerable to dark pattern applications (McStay & Rosner, 2021).

Research with children and parents shows that many are already concerned by the lack of transparency, commercial agendas, and risk of deception associated with smart devices and voice assistants (Keymolen & Van der Hof, 2019). The term "creepy" appears numerous times in the literature and is used by children and younger children to describe what they see as the unsettling or frightening aspects of AI technologies (Kucirkova & Hiniker, 2023; Rubegni et al., 2022; Garg & Sengupta, 2020). In one study, children said it was creepy and misleading for a Roomba vacuum to "talk" to them using their parent's voice (recordings) (Yip et al., 2019).

2.2.2 LEARNING

Data-centric technologies are widely used in elementary and high schools, as well as early childhood education settings (Bradbury & Roberts-Holmes, 2018). Globally, educational technologies ("EdTech") fuel a large market sector and its products increasingly feature AI (Tobin, 2023)[2].Overall, there is a lot of optimism about AI's potential to benefit both learners and educators moving forward. The literature shows that under the right conditions, AI-driven tools can enhance learning for children, including the very young (Kewalramani et al., 2021; Lin 2022).

As Druga et al. (2023) describe, AI systems can help children by improving online search quality, voice assistants can bolster children's access to information, and tutoring chatbots can provide personalized feedback and learning experiences. Concurrently, AI can help educators track, evaluate, and personalize student learning, while automating onerous administrative tasks (Cardona et al., 2023). The literature shows that building critical data and algorithmic literacies among both groups (students and teachers) is crucial for realizing these potential benefits (Ciccone, 2023).

However, researchers warn that the emphasis on efficiency found in AI-driven tools can flatten out differences among students and undermine the tailoring of curricula to individual students that teachers already engage in (Selwyn, 2019). There is a lack of oversight in how EdTech, including AI, is used in schools and early childhood education settings. Meanwhile, massive amounts of student data are collected at and by schools, the combined result of government mandates, a data-centric EdTech ecosystem, and school policies (Livingstone & Pothong, 2022).

Similar trends are found in cultural institutions (e.g., public libraries), child welfare and protection services, and hospital and medical services, where data-centric technologies are used to automate administrative tasks and to track and classify children increasingly involve AI (Hoodbhoy et al., 2021; Saxena et al., 2020).

As privately-owned AI technologies spread across public education systems worldwide, there is a growing need for critical research on their designs, data collection practices, and impacts. There is also an urgent need to build critical AI literacy among children from kindergarten to grade 12 and beyond (UNESCO, 2022). The digital literacies of Canadian children are low, uneven, and largely correlated with parental literacy levels and practices (Donelle et al., 2021). Children's and parents' access to literacy supports (e.g., curriculum, training) varies wildly across age, socio-economic, race and other demographic categories.

The literature emphasizes the benefits of hands-on making of content and code for children's digital literacies, especially when they have opportunities to share their creations, collaborate and receive feedback from others (Holbert et al., 2020; Lankshear & Knobel, 2011; Fields & Grimes, 2020). These findings are consistent with established methods for supporting media and textual literacy acquisition (e.g., Buckingham, 2019).

Lastly, children, young children, and adolescents engage in significant amounts of "informal learning" outside of school (Gee, 2007). Under the right conditions, young people can develop multiple literacies (critical, computing, algorithmic, etc.) by playing, consuming, and interacting with data-centric technologies at home and in other out-of-school contexts (Jenkins, 2009; Dasgupta & Hill, 2023). Supporting such opportunities is especially important for minoritized children and adolescents living in under-resourced communities (Pinkard, 2019).

[2] According to data analyst firm Global Data, the Canadian EdTech sector alone generated $1.9 billion in 2022 https://www.globaldata.com/store/report/canada-edtech-market-analysis/.

2.3

SOCIAL/PARASOCIAL RELATIONSHIPS

Using data-centric technologies often involves interacting with other people. This can be a valuable source of social support, creative collaboration, and civic and community engagement. For example, during the pandemic, many young children used videoconferencing to stay in touch with grandparents (Côté et al., 2022). Similarly, some children and adolescents with disabilities use data-centric technologies to build meaningful relationships with peers (e.g., Alper, 2023). Often, the people that young users engage with using data-centric technologies are incredibly important to them (Ito et al., 2010).

The social dimensions of data-centric technologies become problematic when other users say or do things that are violent, hateful, or otherwise harmful. For example, one study showed that 25% of Canadians aged 12 to 17 years had experienced cyberbullying in the past year (Statistics Canada, 2022). Rates were significantly higher (52%) among non-binary youth. Research conducted in the US shows similarly elevated risks for BIPOC children, specifically Black youth, many of whom experience online racial discrimination multiple times daily (English et al., 2020). While some scholars are optimistic that AI can reduce exposure to harmful content and people while increasing the efficiency and transparency of content moderation systems (e.g., Singh et al., 2022), the research also shows that AI can exhibit bias and amplify discrimination (Nahmias & Perel, 2021; Siapera, 2021).

The literature indicates that children and adolescents primarily think about their online interactions within the context of social relationships (e.g., Stoilova et al., 2019). Notably, their interactions with data-centric technologies are often enmeshed in existing social relationships. For some, this starts before they are even born, as data is created and shared about them through parents' pregnancy apps and social media posts (Barassi, 2020). In many households, smart home devices record, track, and learn children's sounds and movements (Neville & Coulter 2022). It is not yet known how long or how far-reaching these digital traces follow children as they age. However, the high risk of harm (e.g., to children's future prospects) led to the inclusion of a "right to be forgotten" in the EU's General Data Protection Rule (GDPR) (Bunn, 2019).

Data-centric technologies are sometimes used to engage in "para-social relationships" (PSR)--one-sided emotional attachments people sometimes develop toward media characters, influencers, or celebrities (Boerman & Reihmersdal, 2020). Research on children's PSR with voice assistants and AI-driven toys reveals potential benefits (e.g., Kewalramani et al., 2021) and risks (e.g., Le et al., 2022). However, concerns that children's relationships with AI will displace human connections are not supported by evidence. Instead, the literature shows that children's feelings about robots, smart toys, and voice assistants are nuanced and distinctive (Kahn et al., 2013; Kory-Westlund et al., 2018; Aguiar, 2021).

2.4

DEVELOPMENT AND WELL-BEING

There is a large body of literature examining the impacts of data-centric technologies on child and adolescent development, health (physical and mental), and well-being. Much of this research has been tentative and unable to establish causality, instead providing evidence of correlation between certain uses of specific data-centric technologies and potential risks or potential benefits, with multiple variables and variations involved (e.g., Hancock et al., 2022).

There is, however, compelling evidence of harm associated with the online migration of interactions and materials already established as harmful in the offline world, such as racial discrimination and child sexual abuse (English et al., 2020; Ringrose & Regehr, 2023). Scholars are concerned that AI will substantially increase the volume and ease with which these harmful materials and interactions are generated and spread (e.g., Karasavva & Noordbhai, 2021).

Concurrently, the research suggests that certain uses of data-centric technologies can be beneficial for young people's emotional development and well-being. For example, several scholars argue that virtual reality can be used to support the development of emotion regulation skills and treat anxiety among adolescents (e.g., Hugh-Jones et al, 2023). Studies conducted with children and adolescents show that most young people believe that having access to positive digital experiences increases their well-being.

Notably, most children and adolescents also believe that using data-centric technologies can at times negatively impact their mental health and safety (Third et al., 2021). Here, young users are most concerned about specific types of interactions, content, and business practices, rather than overall "screen time" effects. For example, many of the 8-to-18-year-olds who took part in the children's consultation for the drafting of the GC25 reported feeling pressured to curate their online identity (Third et al., 2021).

Meanwhile, common business practices such as persuasive design fail to account for children's developing capacities and inexperience. In some cases, children's lack of knowledge is actively exploited for commercial gain (e.g., Staba & Moore, 2023). Children's understanding of how data-centric technologies work develops over time, but even most adolescents display a relatively poor understanding of data collection practices (Stoilova et al., 2019). A recent study of Australian 12-to-16-year-olds found that nearly half (47%) had never heard the term algorithm associated with online news (Notley et al., 2023). It is questionable that children of any age can give truly informed consent to the complex and ambiguous processes driving AI.

2.4

PRIVACY AND OTHER RIGHTS

Children's rights in the digital environment are of increasing public and academic interest. For example, recent works examine how children's right to play is supported in technology design and policy (e.g., Livingstone & Pothong, 2022), and how children's right to participate in decisions that impact them confirms the need for responsible, child-inclusive AI research and development (e.g., Ito et al., 2023). Overall, the literature is heavily focused on children's privacy rights.

Data-centric tech companies have a long history of infringing on the privacy rights of children, young children, and adolescents. The literature shows that young users' data, including their interactions with friends and family, are frequently collected and used to build detailed profiles about them or to influence their beliefs and behaviours (Turow, 2021; Steeves, 2016; Srivastava et al., 2023). Meanwhile, sensitive student data is collected by EdTech companies with a history of data and privacy breaches (Selwyn et al., 2020).

Scholars highlight the unfairness of placing the onus on children, adolescents, and parents/caregivers to know and manage the complex, often obscured, impacts that data-centric technologies have on children's privacy and other rights (e.g., Takhshid, 2023). Positioning this as a matter of individual consumer choice ignores the high social and economic costs associated with non-use (e.g., D'Lima & Higgins, 2021), and the literacy deficits found among both children and parents (e.g., Vittrup et al., 2014). It also fails to address the systemic biases that data-centric technologies often reflect and reproduce (Benjamin, 2019; O'Neill et al., 2022).

Scholars call attention to the fact that children's data is often collected as part of family data, classroom data, multi-user and "public" data sets. For example, vast amounts of data are passively gathered from children in homes and at schools by smart devices and monitoring software (Phippen & Brennan, 2020). Children's agency and informed consent are largely omitted in these contexts. Meanwhile, parents are tasked with managing (e.g., setting parental controls), assessing, and consenting to a prolific and ever-growing number of companies, applications, and devices that collect children's data (Barassi, 2020). Some parents share data or intimate details about their children online without their consent (Plunkett, 2020).

Scholars argue that the risks to children's privacy and other rights in the digital environment are heightened by AI. For instance, children's lack of emotional maturity and data literacy, combined with AI-driven micro-targeting and "emotional AI," increases the risk of commercial exploitation (Van der Hof et al., 2020). Meanwhile, the emerging market for AI-driven age verification recalls research showing that these and other "child safety" products often infringe on children's privacy, and cultural, and participatory rights (Geist, 2022; Shade, 2011).

Parents have a central role to play in safeguarding children's privacy and supporting children's agency vis-a-vis AI. However, there are significant disparities in parents' access to the technologies, knowledge, literacy and other resources needed to effectively fulfil this role (Druga et al., 2022).

3.0

A FRAMEWORK FOR RESPONSIBLE AI AND CHILDREN

In reviewing the literature, we identified two recurring themes that warrant immediate attention: the need to identify best practices for involving children in the design and development of AI technologies; and the need to shift regulatory efforts to children's rights rather than focusing too narrowly on privacy as a form of consumer data protection.

3.1

DEVELOPING AI WITH AND FOR CHILDREN

The literature emphasizes that children should not simply be protected but also empowered in their interactions with AI. Less attention has been paid, however, to practical strategies for involving children in technology development. We reviewed fifty studies to evaluate how children of diverse ages have contributed to date to the development and research of emerging AI technologies (Veldhuis et al., 2024). We found that the degree of authenticity and personal engagement varies depending on the methodology and its implementation.

The literature shows that hands-on activities are especially crucial for effectively engaging children of diverse backgrounds and ages. Below is a description of four types of activities known to support children's engagement in both thinking about and creating alternatives for AI technologies. The activity types follow the typical stages of the design process. For each activity, examples from the literature involving specific age groups are provided. However, all four activity types have been effectively implemented across age groups. Additional recommendations for involving children and adolescents in AI research and development are provided in Appendix 1.

3.1.1 SENSITIZING ACTIVITIES

Sensitizing activities typically aim to provoke critical reflection on existing technologies by engaging children of all ages with the ethical implications or workings of the technology.

In most cases, stories are used to elicit reflection on the ethical implications of AI, paired with personal, hands-on exploration.

Sensitizing activities should be used to both gauge and build children's knowledge levels to ensure that their understanding is roughly at the same level and that they have equal opportunities to participate.

EXAMPLES

1. Explorations of bias in Google search engines with 8-to10-year-olds (Irgens et al., 2022);
2. Critical analyses of AI-generated poetry with adolescents aged 15 years and older (Lee et al., 2022).

3.1.2 REFLECTION ACTIVITIES

Insights from sensitizing activities can help children and adolescents reflect on requirements for AI technologies in their personal world. This includes

- Requirements for what the technology should be able to do;
- Situations in which the technology might be helpful or harmful;
- Reflection on the values of stakeholders; and
- Preferred and harmful ways of interaction.

The latter has particular implications for AI technologies, since some interactive AI technologies, such as voice assistants or social robots, might be perceived by children as having personalities.

Children can be prompted to reflect on how the technology might impact others. They should be encouraged to investigate the groups that benefit from AI technologies as well as the groups that may be adversely affected by them.

EXAMPLES

1. Exploring how and whom digital assistants help or harm with 11-to-12-year-olds (Solyst et al., 2022);
2. Inviting 13-to14-year-olds to collaborate with a robot agent to determine the robot's vocabulary and personality (Li et al., 2023).

3.1.3 DESIGN-ORIENTED ACTIVITIES

Children can reflect on the real-world implementation of these requirements and their ethical and societal impacts through design-oriented activities. The main characteristic of such activities is that they allow children to create alternative scenarios.

Constructive design activities, such as prototyping with or without technical materials, can provide children with an opportunity to reflect on how to interact with the technologies they have designed.

Children should also be encouraged to define their design opportunities instead of attempting to solve a problem for which AI might not even be a viable solution.

EXAMPLES

1. Storyboarding with 11-to-12-year-olds (Buddemeyer et al., 2022);
2. Low-fidelity prototyping with 8-to-10-year-olds (Garg & Sengupta, 2020), or with children under 8 years (Mott et al., 2022).

3.1.4 EVALUATION ACTIVITIES

Children can reflect on the related problems in their newly created scenarios through evaluation activities. This can be accomplished by acting out the scenario or interacting with the prototype.

By presenting their designs, children can also solicit feedback from others. Children can then use their new insights to iterate on their design and update their scenarios.

It would be beneficial to implement reflection activities to help children distill how they might apply the insights they have gained to their futures.

EXAMPLES

1. Inviting 13-to-14-year-olds to write a letter to their future selves advising how to interact with AI technologies (Garg, 2021);
2. Creating videos with children under 8 years in which they appear as experts on the technology (Mott et al., 2022).

4.0

RIGHTS-BASED APPROACH TO REGULATION

The privacy debate emerging around children and AI is rooted in older debates about the impacts of data-centric technologies on children's rights and well-being. These debates provide important lessons for policymakers to keep in mind as they turn their attention to AI.

First and foremost, attempts to regulate AI have been mired in the push and pull between the desire to promote innovation and concerns about harm, especially to children. This same push and pull has shaped privacy legislation from the start (Mackinnon & Shade, 2020; Reyes et al., 2018).

When the Canadian government first introduced private sector privacy legislation, the Personal Information Protection and Electronic Documents Act (PIPEDA), it was quickly framed as a trade issue. Although children were not specifically mentioned in PIPEDA, they were touted as natural technology users who would drive economic growth so long as regulation did not unduly burden tech companies (Shade, 2011). The Act was drafted to create the kind of consumer trust that would ensure people, including children, continued to participate in the online economy, with their data fueling innovation (Steeves, 2009). The research shows, however, that PIPEDA has failed to develop consumer trust, especially among Canadian children and adolescents (Micheti et al., 2010; Third et al., 2021).

Child advocates such as UNICEF, the 5Rights Foundation (UK), Child Rights Connect (Switzerland), and the Coalition for the Rights of the Child (Canada) argue that a child rights approach would better protect children and children's interests in their relationships and interactions with AI. One advantage is that the UNCRC addresses several issues, opportunities and challenges implicated in AI, including children's right to privacy, as well as their rights to access information, play and participate in cultural life, and to be free from discrimination, commercial exploitation, and abuse. A rights-based model captures more of children's lived experiences and needs than a data protection model (Steeves, 2023).

An example can be found in the UK's Age Appropriate Design Code. The UK Information Commissioner worked with child advocates and tech corporations to create a set of principles to ensure that platforms would be developed in ways that are respectful of children's rights and best interests from the outset. Although the majority of its 15 clauses do focus on privacy, the code implements child rights language in three notable instances: it makes the best interests of the child a primary consideration for designers; it tells designers not to use children's data in ways that have been shown to be detrimental to their wellbeing; and it only enables profiling if there are measures in place to protect children from harm, especially harm from seeing content that negatively impacts their health or wellbeing.

Blanket prohibitions on the infringement of children's rights are much more likely to constrain the problematic industry practices standing in the way of responsible AI for children. However, reforms to PIPEDA, currently before Parliament, weaken provisions that have historically protected children's privacy (e.g., s. 5(3) see OPC, 2021). To date, legislators have also failed to act on the Privacy Commissioner of Canada's recommendations that the Bill recognize privacy as a fundamental right and enact the best interests of the child as an enforceable standard (OPC, 2023). They have yet to respond to the new requirements set out in the UNCRC GC25.

Whether child rights language can push back against more instrumental approaches to "responsible AI" is yet to be seen. However, any legislative efforts will be measured against the broad commitment to child rights made by UNCRC signatory states, including Canada.

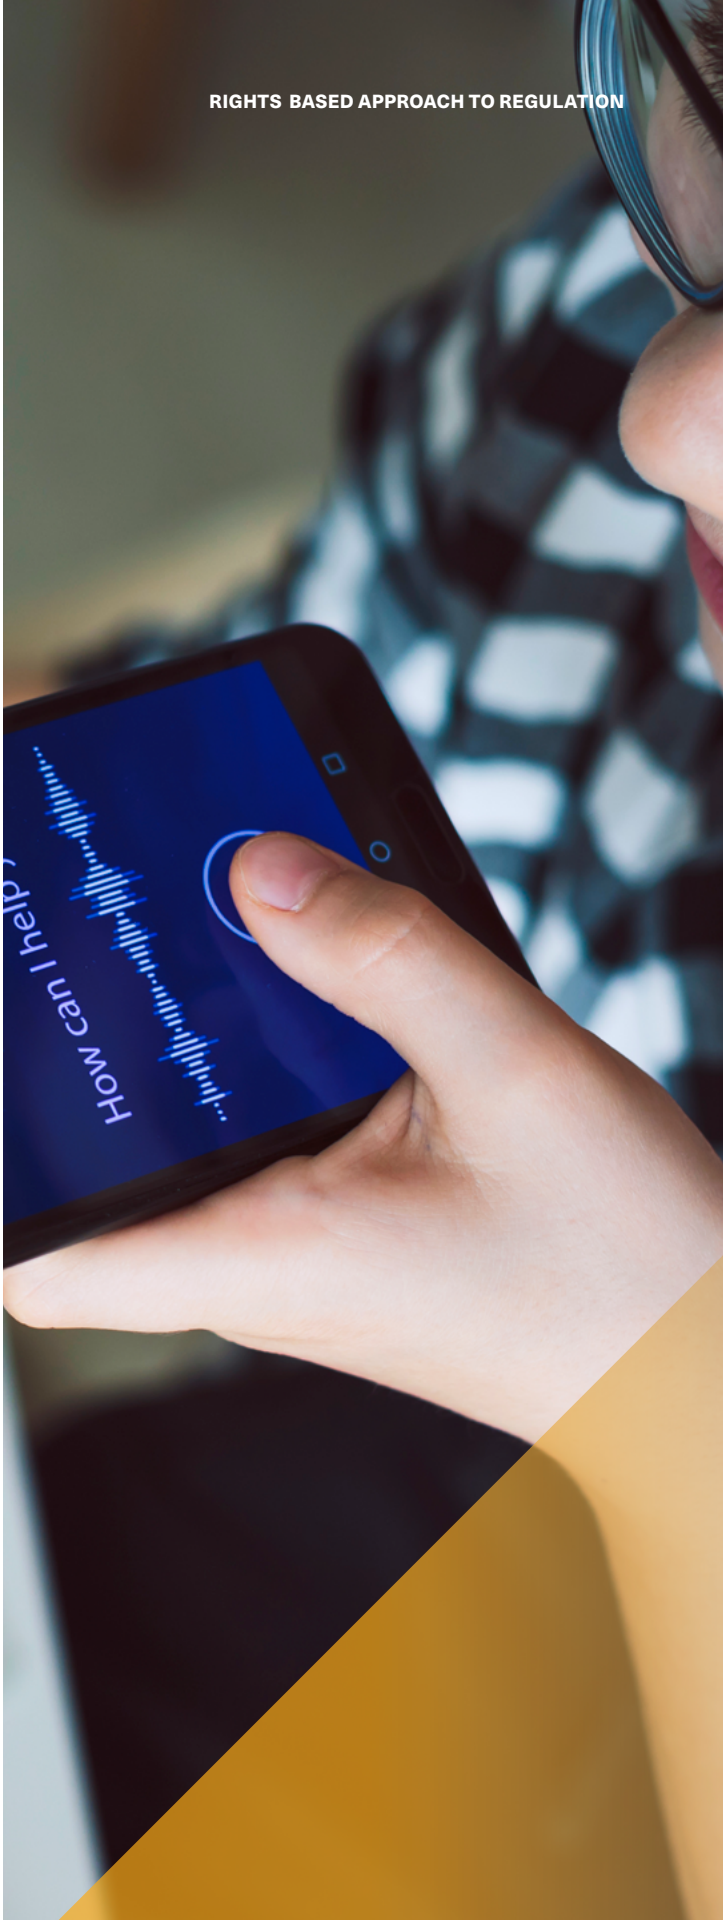

5.0
KEY TAKEAWAYS

1 **WE MUST MOVE PAST QUESTIONS ABOUT CAUSALITY TO CONSIDER CORRELATIONS, POTENTIAL BENEFITS, AND POTENTIAL RISKS WHEN THINKING ABOUT HOW TO BUILD A RESPONSIBLE AI ECOSYSTEM FOR CHILDREN.** This requires looking beyond child development research to include fields examining other crucial dimensions (e.g., social, cultural, educational) of children's relationships with technologies and the tech industries. While more research on children and AI is needed, the "no existing research" argument is misleading.

2 **FUTURE POLICY DISCUSSIONS SHOULD INCLUDE CONSULTATIONS WITH SCHOLARS REPRESENTING THE MULTIPLE AND DIVERSE FIELDS ENGAGED IN THIS RESEARCH.** Most of the leading research and theories about children and data-centric technologies is interdisciplinary, which is reflective of the complex impacts these technologies, industries, and policies have on children's lives.

3 **THE GC25 CAN SERVE AS A GUIDE FOR SETTING PRIORITY AREAS AND FLAGGING CONCERNS** that children, adolescents, experts, caregivers, educators, companies, and child advocates from around the world have already flagged as paramount to supporting children's rights and wellbeing in the digital environment. The GC25 applies to all digital technologies, including AI.

4 **CHILDREN OF ALL AGES SHOULD BE INVOLVED IN RESEARCH, POLICY DECISIONS, TECHNOLOGY DEVELOPMENT AND DESIGNS THAT ARE GOING TO IMPACT THEM.** Efforts should apply a child-centred design methodology and emphasize age-appropriate hands-on activities that enable children to learn about, reflect on, design, and evaluate AI technologies and policies.

5 **ANY DISCUSSIONS ABOUT RESPONSIBLE AI AND CHILDREN MUST CONSIDER THAT THERE IS NO SUCH THING AS A "UNIVERSAL" CHILDHOOD OR YOUTH EXPERIENCE.** Children are an incredibly diverse population whose interactions with technologies are shaped by individual personal, familial, cultural, socio-economic, and geographic contexts.

6 **SUPPORT AND FUNDING FOR NEW, CRITICAL, AND INTERDISCIPLINARY RESEARCH ON CHILDREN AND AI IS IMPERATIVE.** There is a growing need for research that considers the wider socio-cultural and political economic impacts of the infiltration of AI technologies across children's lives—at home and at school, in public and private, online and off.

6.0
REFERENCES

Aguiar, N.M. (2021). A paradigm for assessing adults' and children's concepts of artificially intelligent virtual characters. Human Behavior and Emerging Technologies, 3(4), 618-34.

Alper, M. (2023). Kids across spectrums: Growing up autistic in the digital age. MIT Press.

Antle, A., & Kitson, A., (2021). 1,2,3,4 tell me how to grow more: A position paper on children, design ethics and biowearables. International Journal of Child-Computer Interaction, 30, 100328.

Barassi, V. (2020). Child data citizen: How tech companies are profiling us from before birth. MIT Press.

Benjamin, R. (2019). Race after technology: Abolitionist tools for the New Jim Code. Polity Press.

Boerman S.C., & van Reijmersdal, E.A. (2020). Disclosing influencer marketing on YouTube to children: The moderating role of para-social relationship. Frontiers in Psychology, 10, 3042.

Bradbury, A., & Roberts-Holmes, G. (2018). The datafication of primary and early years education: Playing with numbers. Routledge.

Buckingham, D. (2019). The media education manifesto. Polity Press.

Buddemeyer, A., Nwogu, J., Solyst, J., Walker, E., Nkrumah, T., Ogan, A., Hatley, L., & Stewart, A. (2022). Unwritten magic: Participatory design of AI Dialogue to empower marginalized voices. Proceedings of the ACM Conference on Information Technology for Social Good (pp. 366-372).

Bunn, A. (2019). Children and the 'Right to be Forgotten': what the right to erasure means for European children, and why Australian children should be afforded a similar right. Media International Australia, 170(1): 37-46

Cardona, M.A., Rodríguez, R.J., & Ishmael, K. (2023). Artificial Intelligence and the future of teaching and learning: Insights and recommendations. U.S. Department of Education, Office of Educational Technology.

Ciccone, M. (2023). Algorithmic literacies: K-12 realities and possibilities. In M. Ito, R. Cross, K. Dinaker, & C. Odgers (Eds). Algorithmic rights and protections for children (pp.15-22). MIT Press.

Costello, M., Barrett-Fox, R., Bernatzky, C., Hawdon, J., & Mendes, K. (2020). Predictors of viewing online extremism among America's youth. Youth & Society, 52(5), 710–727.

Côté, I., Gervais, C., Doucet, S., & Lafantaisie, V. (2022). "Je m'ennuie beaucoup de mamie et papi." Impacts des mesures sociosanitaires sur les liens grands-parents et petits-enfants. Enfances, Familles, Générations,(40).

D'Lima, P., & Higgins, A. (2021). Social media engagement and Fear of Missing Out (FOMO) in primary school children. Educational Psychology in Practice, 37(3), 320–338.

Donelle, L., Facca, D., Burke, S., Hiebert, B., Bender, E., & Ling, S. (2021). Exploring Canadian children's social media use, digital literacy, and quality of life: Pilot cross-sectional survey study. JMIR Formative Research, 5(5), e18771–e18771.

Druga, S., Christoph, F., & Ko, A.J. (2022). Family as a third space for AI literacies: How do children and parents learn about AI together? CHI '22, Apr. 29-May 5, New Orleans, LA.

Druga, S., Yip, J., Preston, M., & Dillon, D. (2023) The 4 As: Ask, adapt, author, analyze. In M. Ito, R. Cross, K. Dinaker, & C. Odgers (Eds). Algorithmic rights and protections for children (pp. 193-231). MIT Press.

English, D., Lambert, S. F., Tynes, B. M., Bowleg, L., Zea, M. C., & Howard, L. C. (2020). Daily multidimensional racial discrimination among Black U.S. American adolescents. Journal of Applied Developmental Psychology, 66, 101068–12.

Fields, D.A., & Grimes, S.M. (2020). Shaping learning online for making and sharing children's DIY media. In N. Holbert, M. Berland, & Y.B. Kafai (Eds.) Designing Constructionist Futures: The Art, Theory, and Practice of Learning Designs (pp.255-263). MIT Press.

Garg, R. (2021). Engaging parents and teens in an asynchronous, remote, community-based method for understanding the future of voice technology. Proceedings of the 20th Annual ACM Interaction Design and Children Conference (pp.224-235).

Garg, R., & Sengupta, S. (2020) He is just like me: A study of the long-term use of smart speakers by parents and children. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 4(1), 1-24.

Gee, J.P. (2007). Good video games + good learning: Collected essays on video games, learning and literacy. Peter Lang.

Geist, M. (2022). Age verification requirements for Twitter or website blocking for Reddit? Michael Geist [Blog]. Newstex.

Giddings, S. (2014). Gameworlds: Virtual media and children's everyday play. Bloomsbury.

Grandinetti, J. (2023). Examining embedded apparatuses of AI in Facebook and TikTok. AI & Society, 38 (4), 1273-1286.

Grimes, S.M. (2021). Digital playgrounds: The hidden politics of children's online play spaces, virtual worlds and connected games. University of Toronto Press.

Hancock, J., Liu, S. X., Luo, M., & Mieczkowski, H. (2022). Psychological well-being and social media use: A meta-analysis. SSRN.

Helsper E. J. (2021). The digital disconnect: The social causes and consequences of digital inequalities. SAGE.

Hiniker, A., Wang, A., Tran, J., Zhang, M.R., Radesky, J., Sobel, K., & Hong, S.R. (2021). Can conversational agents change the way children talk to people? IDC '21, June 24-30, Athens, Greece.

Hugh-Jones, S., Ulor, M., Nugent, T., Walshe, S., & Kirk, M. (2023). The potential of virtual reality to support adolescent mental well-being in schools: A UK co-design and proof of concept study. Mental Health & Prevention, 30, 1-7.

Irgens, G.A., Adisa, I., Bailey, C., & Quesada, H. V. (2022). Designing with and for youth. Educational Technology & Society, 25(4), 126-141.

Ito, M., Horst, H., Bittanti, M., & boyd, d. (Eds.) (2010). Hanging out, messing around, geeking out: Living and learning with new media. MIT Press.

Ito, M., Cross, R., Dinaker, K., & Odgers, C. (2023). Algorithmic rights and protections for children. MIT Press.

Jenkins, H. (2009). Confronting the challenges of participatory culture: Media education for the 21st century. MIT Press.

Kafai, Y. B., & Fields, D.A. (2013). Connected play: Tweens in a virtual world. MIT Press.

Kahn, P. H., Gary, H. E., & Shen, S. (2013). Children's social relationships with current and near-future robots. Child Development Perspectives, 7, 32–37.

Karasavva, V., & Noorbhai, A. (2021). The real threat of deepfake pornography: A review of Canadian policy. Cyberpsychology, Behavior, and Social Networking, 24(3), 203-209.

Kewalramani, S., Kidman, G., & Palaiologou, I. (2021) Using Artificial Intelligence (AI)-interfaced robotic toys in early childhood settings: A case for children's inquiry literacy. European Early Childhood Education Research Journal, 29(5), 652-668.

Keymolen, E., & Van der Hof, S. (2019). Can I still trust you, my dear doll? A philosophical and legal exploration of smart toys and trust. Journal of Cyber Policy, 4(2), 143–159.

Konstantoni, K., & Emejulu, A. (2017). When intersectionality met childhood studies: The dilemmas of a travelling concept. Children's Geographies, 15(1), 6-22.

Kory-Westlund, J. M., Park, H. W., Williams, R., and Breazeal, C. (2018). Measuring young children's long-term relationships with social robots. Proceedings of the 17th ACM Conference on Interaction Design and Children (pp.207–218).

Kucirkova, N., & Hiniker, A. (2023). Parents' ontological beliefs regarding the use of conversational agents at home: Resisting the neoliberal discourse. Learning, Media and Technology.

Lankshear, C., & Knobel, M. (2011). New literacies: everyday practices and social learning. Open University Press.

Lee, C. H., Gobir, N., Gurn, A., & Soep, E. (2022). In the black mirror: Youth investigations into artificial intelligence. ACM Transactions on Computing Education, 22(3), 1-25.

Li, Y., Nwogu, J., Buddemeyer, A., Solyst, J., Lee, J., Walker, E., Ogan, A., & Stewart, A.E. (2023). "I want to be unique from other robots": Positioning girls as co-creators of social robots in culturally-responsive computing education. Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (pp. 1-14).

Livingstone, S., & Pothong, K. (2022). Imaginative play in digital environments: Designing social and creative opportunities for identity formation. Information, Communication & Society, 25(4), 485-501.

Mackinnon, K., & Shade, L.R. (2020). "God only knows what it's doing to our children's brains": a closer look at internet addiction discourse. Jeunesse: Young People, Texts, Cultures 12(1), 16-38.

Mascheroni, G., & Holloway, D. (2019). The Internet of toys: Practices, affordances and the political economy of children's smart play. Springer.

Mascheroni, G., & Siibak, A. (2021). Datafied childhoods: Data practices and imaginaries in children's lives. Peter Lang.

McEwen, R. (2023). AI is not intelligent and needs regulation now (OpEd). The Hill Times, September 27.

McStay, A., & Rosner, G. (2021). Emotional artificial intelligence in children's toys and devices: Ethics, governance and practical remedies. Big Data & Society, (Jan-Jun), 1-16.

MediaSmarts. (2022). Young Canadians in a Wireless World, Phase IV: Life Online.

Micheti, A., Burkell, J., & Steeves, V. (2010). Fixing broken doors: Strategies for drafting privacy policies young people can understand. Bulletin of Science, Technology & Society, 30, 130-143.

Mott, T., Bejarano, A., & Williams, T. (2022, March). Robot co-design can help us engage child stakeholders in ethical reflection. Proceedings of the 17th ACM/IEEE International Conference on Human-Robot Interaction (pp. 14-23).

Muralidharan, V., Burgart, A., Daneshjou, R., & Rose, S. (2023). Recommendations for the use of pediatric data in artificial intelligence and machine learning ACCEPT-AI. Digital Medicine 6, 166.

Nahmias, Y., & Perel, M. (2021). The oversight of content moderation by AI: Impact assessments and their limitations. Harvard Journal on Legislation, 58(1), 145-194.

Neville, S., & Coulter, N. (2022). Sound the alarm on Alexa and the eavesmining industry. Communications Law, 27(4), 194-197.

Noble, S. (2018). Algorithms of oppression: How search engines reinforce racism. NYU Press.

Notley, T., Chambers, S., Zhong, H.F., Park, S., Lee, J & Dezuanni, M. (2023) News and Young Australians in 2023: How children and teens access, perceive and are affected by news media. Research Report, Western Sydney University.

O'Neill, C., Selwyn, N., Smith, G., Andrejevic, M., & Gu, X. (2022). The two faces of the child in facial recognition industry discourse: biometric capture between innocence and recalcitrance. Information, Communication, & Society, 25(6), 752-767.

Office of the Privacy Commissioner (OPC) (2023). OPC's 15 key recommendations on Bill C-27. News Release, May 11.

OPC (2021). Policy position of online behavioural advertising. Policy Position, August 13.

Pangrazio, L., & Mavoa, J. (2023). Studying the datafication of Australian childhoods: Learning from a survey of digital technologies in homes with young children. Media International Australia, 0(0), 1-16.

Pinkard, N. (2019). Freedom of movement: Defining, researching, and designing the components of a healthy learning ecosystem. Human Development, 62, 40-65.

Plunkett, L. (2019). Sharenthood: How the digital tech habits of parents, teachers, and other trusted adults harm kids and teens. MIT Press.

Radesky, J., Hiniker, A., McLaren, C., Akgun, E., Schaller, A., Weeks, H.M., Campbell, S., & Gearhardt, A.N. (2022). Prevalence and characteristics of manipulative design in mobile applications used by children. JAMA Network Open - Pediatrics, 5(6), 1-11.

Reyes, I., Wijesekera, P., Reardon, J., Elazari Bar On, A., Razaghpanah, A. Vallina-Rodriguez, N., & Egelman, S. (2018). "Won't somebody think of the children?" Examining COPPA compliance at scale. Proceedings on Privacy Enhancing Technologies (3), 63–83.

Ringrose, J., & Regehr, K. (2023). Recognizing and addressing how gender shapes young people's experiences of image-based sexual harassment and abuse in educational settings. Journal of Social Issues.

Rubegni, E., Malinverni, L., & Yip, J.C. (2022). "Don't let the robots walk our dogs, but it's ok for them to do our homework": Children's perceptions, fears, and hopes in social robots. IDC '22, June 27–30, Braga, Portugal.

Saxena, D., Badillo-Urquiola, K., Wisniewski, P., & Guha, S. (2020). A human-centered review of the algorithms used within the US child welfare system. CHI'20, April 25–30, Honolulu, HI.

Siapera, E. (2021). AI content moderation, racism and (de)coloniality. International Journal of Bullying Prevention, 4, 55-65,

Singh, S., Greaves, D.J., & Epiphaniou, G. (2022) A framework for integrating responsible AI into social media platforms. Competitive Advantage in the Digital Economy (CADE 2022) Hybrid Conference, Venice, Italy, 13-15 Jun, 117-120.

Selwyn, N. (2019). Should robots replace teachers? AI and the future of education. Polity Press.

Selwyn, N., Hillman, T., Eynon, R., Ferreira, G., Knox, J., Macgilchrist, F., & Sancho-Gil, J.M. (2020). What's next for Ed-Tech? Critical hopes and concerns for the 2020s. Learning, Media and Technology, 45(1), 1-6.

Shade, L.R. (2011). Surveilling the girl via third and networked screens. In Kearney, M.C. (Ed.) Mediated girlhoods: New explorations of girls' media culture [pp. 261-276]. Peter Lang.

Solyst, J., Axon, A., Stewart, A. E., Eslami, M., & Ogan, A. (2022). Investigating girls' perspectives and knowledge gaps on ethics and fairness in Artificial Intelligence in a lightweight workshop. arXiv preprint arXiv:2302.13947.

Srivastava, S., Wilska, T.-A., & Nyrhinen, J. (2023). Awareness of digital commercial profiling among adolescents in Finland and their perspectives on online targeted advertisements. Journal of Children and Media, 17(4), 559–578.

Staba, K., & Moore, D. (2023). Over a Fortnite: Epic Games to pay record US$520 million to settle FTC claims of children's privacy violations and digital dark patterns that continued for years. Computer and Internet Lawyer, 40(4), 11–13.

Steeves, V. (2023). Children's online privacy: Uses, abuses and rights. In Harrison, V., Collier, A., & Adelsheim, S. (Eds.) Social media and youth mental health: A public health perspective. American Psychiatric Association.

Steeves, V. (2016). Swimming in the fishbowl: Young people, identity and surveillance in networked spaces. In van der Ploeg, I., & Pridmore, J. (Eds.), Digitizing Identities. Routledge.

Steeves, V. (2009). Data protection versus privacy: Lessons from Facebook's Beacon. In Matheson, D. (Ed.) The Contours of Privacy [pp. 183-196]. Cambridge Scholars Press.

Stevens, R., Gilliard-Matthews, S., Dunaev, J., Todhunter-Reid, A., Brawner, B., Stewart, & Stewart, J. (2017). Social media use and sexual risk reduction behavior among minority youth: Seeking safe sex information. Nursing Research, 66(5), 368-377

Stoilova, M., Livingstone, S., & Nandagiri, R. (2019). Children's data and privacy online:

Growing up in a digital age. London School of Economics and Political Science.

Takhshid, Z. (2023). Children's digital privacy and the case against parental consent. Texas Law Review, 101(6), 1417-1455.

Third, A., & Moody, L. (2021). Our rights in the digital world: A report on the children's consultations to inform UNCRC General Comment 25. 5Rights Foundation.

Tobin, J. (2023). Educational technology: Digital innovation and AI in schools. In Focus (Nov. 20). House of Lords Library.

Turow, J. (2021). The voice catchers: How marketers listen in to exploit your feelings, your privacy, and your wallet. Yale University Press.

UNESCO (2022). K-12 AI curricula: a mapping of government-endorsed AI curricula. https://unesdoc.unesco.org/ark:/48223/pf0000380602

Van der Hof, S., Lievens E, Milkaite I, Verdoodt V, Hannema T, & Liefaard T. (2020) The child's right to protection against economic exploitation in the digital world. International Journal of Child Rights, 28, 833–859.

Veldhuis, A., Kenny, S., Lo, P., & Antle, A.N. (2024). Existing literature on child-centered design and the ways children have been involved in AI research and development. Report. School of Interactive Arts and Technology, Simon Fraser University.

Vittrup, B., Snider, S., Rose, K.R., & Rippy, J. (2014). Parental perception of the role of media and technology in their young children's lives. Journal of Early Childhood Research, 14(1), 43-54.

Willis, L.E. (2020). Deception by design. Harvard Journal of Law & Technology, 34(1), 115-190.

Wronski, L. (2019). Young children and smart speakers. Common Sense Media/Survey Monkey.

Yip, J.C., Sobel, K., Gao, X., Hishikawa, A.M., Lim, A., Meng, L., Ofiana, R.F., Park, J, & Hiniker, A. (2019). Laughing is scary, but farting is cute: A conceptual model of children's perspectives of creepy technologies. CHI 2019, May 4-9, Glasgow, UK.

APPENDIX 1

RECOMMENDATIONS FOR INVOLVING CHILDREN AND ADOLESCENTS IN AI RESEARCH AND DEVELOPMENT

DETERMINE EXISTING AI LITERACY LEVELS:

- Use short provocations (e.g., quizzes) to create a baseline proficiency for subsequent activities.
- Break down common and naive notions of AI technologies from popular culture to provide realistic understandings of AI functionalities and limitations.

ACTIVE ENGAGEMENT THROUGH DESIGN PROCESSES:

- Alternate activities in which children engage in the ideation and design, with critical thinking activities in which they analyze and reflect.
- Inform children about the design process itself and the need for iteration, research, and continuous reflection.

MULTIPLE AND INTERRELATED SENSEMAKING ACTIVITIES:

- Have a series of diverse and interrelated activities that reinforce AI concepts from different angles.
- Provide experiences and tools to support comprehensive understanding and elicit meaningful contributions from children.

REFLECTION ON BROADER CONTEXTUAL IMPLICATIONS:

- Strengthen the meaning and authenticity of children's voices by fostering their critical thinking about AI's impact on society and its broader socio-political contexts.
- Help children identify and voice ethical concerns about AI implementation through deconstruction exercises and Value Sensitive Design (VSD) methods.

HOLISTIC AND MULTIFACETED PERSPECTIVES:

- Help children consider ethical consequences as well as the possibilities of AI systems.
- Encourage a diverse range of perspectives (interdisciplinary and intersectional) to develop well-rounded thinking about both the harms and benefits of AI.

MOTIVATION AND TAILORED IMPLEMENTATION:

- Motivate children to participate by tailoring and implementing activities with topics or goals that are relevant or interesting to them. Center activities around personal engagement.
- Where possible, give children the ability to actualize design ideas into high fidelity prototypes.

FOSTER COMMUNITY AND REDUCE POWER IMBALANCES:

- Dedicate time to developing rapport and long-term relationships with children to reduce power dynamics and encourage authentic discussions of AI.
- Community and safe spaces are particularly important when working with marginalized/minority groups or when dealing with sensitive topics.

CIFAR

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

www.cifar.ca/ai