

Project Workshop Report

Generation AI

Establishing Global Standards for Children and AI

June 2019



World Economic Forum
91-93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland
Tel.: +41 (0)22 869 1212
Fax: +41 (0)22 786 2744
Email: contact@weforum.org
www.weforum.org

© 2019 World Economic Forum. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, or by any information storage and retrieval system.

Contents

Background and context	4
Discussion	5
Child rights	6
Plenary overview: Flash talks	6
Breakout group discussions	6
Privacy	6
Algorithms for children	7
Agency	8
Corporate governance	9
Plenary overview: Moderated panel conversation	9
Breakout group discussions	9
Internal processes	9
Public education	10
Consumer protection	10
Assessment and evaluation	11
Public policy	12
Plenary overview: Panel presentations	12
Breakout group discussions	12
Laws and regulation protecting children	12
Government protection	13
Science to policy	14
Policy guidance for AI and child rights	14
Outcomes	16
Workshop participants	17
World Economic Forum contact	17

Background and context

Artificial intelligence (AI) carries with it the promise of enhancing human potential and improving upon social outcomes where existing systems have fallen short. Numerous risks and uncertainties, however, must be addressed as AI continues to evolve and integrate into public and private decision-making systems that define the world and, in particular, the world of opportunity for the people born to it. As digital natives, perhaps no group will be more affected by AI than children. It thus warrants special care to ensure that it is built to uphold children's rights and maximize their developmental growth.

On 6-7 May 2019, the World Economic Forum Centre for the Fourth Industrial Revolution and its partners UNICEF and the Canadian Institute for Advanced Research (CIFAR) hosted a workshop in San Francisco on the joint "Generation AI" initiative. Comprised of key stakeholders from business, academia, government and civil society, the Generation AI community is committed to driving multistakeholder policy solutions that enable opportunities of AI for children while minimizing its potential harms.

This workshop identified deliverables in two key areas:

1. A set of public policy guidelines that direct countries on creating new laws focused on children and AI
2. A corporate governance charter that guides companies leveraging AI to design their products and services with children in mind.

To inform the development of these two deliverables, the workshop was divided into three main sections: child rights, corporate governance and public policy.

Each section began with a plenary overview canvassing relevant issues and leadership perspectives from the variety of experts in the room. Following the plenary, the workshop participants divided into breakout groups focused on a more granular subset of issues under the session's main rubric.

The goal at each stage was to consider the issues through an interdisciplinary lens, evaluating cutting-edge perspectives from media experts and developmental psychology, for example, alongside insight from business and legal practitioners. The result was a rich foundation of material on which the Forum and its Generation AI partners can begin to tangibly structure the governance mechanisms that the project has identified as current priorities.

Discussion



This section captures the conversation that took place during each of the three main sections: child rights, corporate governance and public policy.

Child rights

Plenary overview: Flash talks

Short presentations or “flash talks” from participants prompted the group to think about child rights. A summary of the key issues discussed during these flash talks in advance of the issue-specific breakout group sessions follows.

Flash talk presenters:

Ronald Dahl, Director, Institute of Human Development, University of California, Berkeley

Mizuko Ito, Director, Connected Learning Lab, University of California, Irvine

Erica Kochi, Co-Founder, UNICEF Innovation, United Nations Children’s Fund, New York

Leveraging developmental windows of opportunity

When it comes to children, no one-size-fits-all policy is possible because their vulnerabilities and opportunities for growth and development vary at different ages. Using a developmental science model allows delving deeper into how different identifying features of each developmental stage can be leveraged to inform positive outcomes with respect to AI-enabled toys and products.

The transition from childhood to adolescence is particularly defined by notable changes in emotional response levels and motivational goals. Understanding these changes can help to identify the natural attractors for adolescent learning, which can be leveraged to inform design and policy measures governing AI that impacts this group. Further, due to the intensity of changes that happen at the adolescent stage, adolescence offers a crucial window of opportunity to influence the development of children in their second decade of life. The potential for the positive impact of AI that has been mindfully designed to support and nurture growth is thus heightened during adolescence and should be evaluated with care.

AI amplifies both risks and opportunities; looking forward, it is beneficial to strive to leverage the unique opportunities at each stage of childhood development to inform AI policy and design. This proposed approach can be contrasted with that taken by the EU in its General Data Protection Regulation (GDPR), which considered issues concerning children’s data from a singular perspective. This frames the reality that sophisticated thinking on developmental science is not currently in the lexicon of most policy-makers, as well as the importance of finding ways to translate these concepts for decision-makers.

Contextualizing UNICEF and Generation AI

Certain AI applications that could have a positive impact, such as using facial recognition technology to evaluate whether a child is malnourished, cannot move forward in the context of UNICEF’s focus on ending child hunger because of issues concerning privacy. This highlights a larger theme of the workshop on the priority of privacy rights in any given situation concerning children and AI, when considered against other rights (in this instance, the right to health).

UNICEF is advancing its work in this area by developing AI literacy among its experts and field practitioners. It is working with states to promote the equitable representation of children in data sets that are being used to train AI that will impact them, and developing policy guidance for countries and companies to ensure that child rights are protected. UNICEF is focused on developing AI and child rights policy guidance for national policy-makers, corporations and the UN system, to help put child rights on the policy agenda.

Influencing cultural norms

The intentions of those who are designing technology are not the same as those using it. Regulating technology will not address the underlying issues that contribute to behavioural problems stemming from culture and society. To address these issues, consideration should be given to influencing cultural norm setting and designing community-based interventions. Such interventions are distributed and difficult to control, but necessary from the perspective of meeting people and their circumstances.

Breakout group discussions

Privacy

Key points:

Balancing trade-offs. Context is key when evaluating privacy issues, and a binary, all-or-nothing approach will no longer work. When are other interests considered a priority? It is important to be aware of both sides to effectively consider the trade-offs.

- Large and unbiased data sets are needed to train fair algorithms and create actionable insight to address major social challenges. It is indispensable to meet the challenge of robust data collection while managing privacy concerns.
- Collecting data from children may in some cases enhance development, but in others it may support structures that oppress their potential for growth and their ability to thrive later in life.

- Should companies consider children’s data sensitive by default?
- How can the benefits of behavioural analytics be captured without chilling the freedom of expression in children?

Considering parents and data privacy. Parents present a third actor in the relationship with data. Considering the role of parents makes cases involving children categorically different from the rest of the discussion on data ownership and privacy.

- It may not always be in the interest of children for their parents to have access to their raw data past a certain age or developmental stage.
- Toys that speak with children might hear reports of abuse or other situations that are harmful to the child. This introduces the question of when a toy has the duty to report potential harm, which raises issues on data collection and surveillance practices that would support such a reporting structure in accordance with privacy goals.
- Should a parent be allowed to sell a child’s data and, if so, at what age should the child recover such agency? How should the value be kept in trust? The exposure that children would face in such a situation is likely to be divided along lines of privilege and parental engagement, which raises concerns about equality under a framework for child rights.

Limiting exposure to commercial advertising. What are the responsibilities of technology platforms regarding privacy rights?

- Advertisers can work around rules to deliver ads to children through online platforms. For example, although child-focused platforms limit forms of paid advertising, they still support entire channels devoted to brands, which could be considered a hypercommercialized form of advertising masquerading as content.
- Tech companies are primarily incentivized to make money from advertising; what mechanisms can be created to prioritize the interests of children, given the dominance of market-based incentives?
- To create momentum for policy, it could be useful to define the harms that advocates are attempting to protect against by limiting child exposure to advertising.

Algorithms for children

Key points:

Optimizing algorithms for learning. One opportunity is to focus on building algorithms that are optimized for learning objectives and steer users in positive ways towards pro-social outcomes, much like books, curricula and other legacy forms of educational programming.

- Any situation in which a child is directed towards an outcome raises issues of algorithmic manipulation and agency. However, implementing algorithmic “nudges” towards long-term goals over short-term motivations can work to enhance autonomy. The key issue here is to identify and promote healthy goals for algorithmic exposure when children are involved.
- Teachers and other adults need to entice children a bit in the short term to engage in behaviour that necessitates the development of creativity, character and critical learning skills. The group discussed the viability of applying this concept to purposeful algorithmic design.

Modelling algorithms after learning. How can better algorithms be created by modelling them after how children learn? The discussion focused on framing how algorithms are limited with respect to certain phases and aspects of learning.

- Children demonstrate certain qualities of learning and absorbing the world, especially curiosity and exploration. These qualities could be used to build more effective algorithms if it were possible to determine how to recreate these processes through AI.
- Children are data-efficient, extracting and processing the most relevant data to learn from just a few examples, whereas AI needs to be fed millions of examples to render accurate judgements.
- The challenges in this area are not entirely technical. One way to support a child-like learning potential for algorithms is to create more awareness of learning models (including developmental vulnerabilities and opportunities), which can filter up to influence design.
- The development of learning models that integrate the “salience features” that children demonstrate would also help to mitigate data concerns.

Introducing new market forces. In the current market system, it is unclear how to meaningfully introduce products for children that are optimized for responsibility over revenue. Is it possible to work through markets and policy to transcend the dominance of big tech in this area?

- In the history of television broadcasting, one might make an analogy to the introduction of public broadcasting, public funded programming that aims to be fun, engaging, development-oriented and educational. Is it possible to create a “public option” for digital technology that can motivate youngsters as users?
- Establishing the social expectations of companies (through media and other channels) helps to create the business case for responsibility in algorithmic product design. This approach does not strive to work beyond commercial objectives but rather within them.
- If the notion that companies ultimately cannot look beyond profit is accepted, the focus should be on elevating new movers, such as socially focused companies, media voices and non-governmental organizations, into the space that can operate from a more neutral perspective.

Remedy. How can a process of remedy be supported when algorithms get it wrong and make decisions that have a negative impact? This question, and the important issue of remedy, requires further consideration.

Agency

Key points:

Championing children. The group considered different forms of authority and implementation that ultimately framed the importance of bringing children into the policy-making process by accounting for their views on the policies that affect them.

- No one stakeholder in the life of children, whether it be business, government or parents, can be relied upon to entirely represent the best interests of the child. Thus capturing and incorporating the perspectives of children are needed.
- Are there ways that certain stakeholders can ensure that a child’s agency is protected against potential abuse of authority or illicit control by another? This question becomes particularly difficult when children do not have present or reliable and loving guardians.
- What responsibility do government and companies have to protect children’s long-term interests against oppressive guardianship? This issue is not new but AI could have new implications for perpetuating inequality and limiting opportunities to develop and thrive.

Considering parents and data agency. To what extent should parents and guardians have control over their children’s data? The subject of how the relationship between parents and their children’s data should be structured was also an important point of discussion in the “Privacy” section, as noted above.

- Parental data access and sharing should be considered through a developmental lens: it may be important from a fairness and dignity perspective to re-establish data agency for children when they begin adolescence, due to the psychological, behavioural and developmental transitions that are taking place at this time.
- Looking at how agency between parents and teenagers has been handled in other contexts would help to inform the type of data agency practices that would best serve young people in the age of AI.
- Policy might establish a data brokerage or the terms for private-run data brokerage platforms for children. A data brokerage is a central entity that would set the terms for data usage on behalf of a group of people. Data powers billion-dollar companies and in the aggregate provides the fuel for automated systems. Data brokerages could allow individuals to participate in their commercial value and help to promote autonomy.

Incorporating children as stakeholders. Young people’s capacity to innovate in response to the challenges and opportunities of the platform economy should be used to help design policy.

- Young people are good at learning from and responding to feedback from online media systems.
- Companies that are designing either explicitly for children or that manage systems that fundamentally affect how they live, develop and process the world should involve them at the board and design levels.
- Children are already co-creators of the digital environment; finding more systematic and meaningful ways to solicit and apply their feedback would be beneficial.

Corporate governance

Plenary overview: Moderated panel conversation

This session's panel was comprised of experts representing business, academia and civil society. Its goal was to examine ideas and priorities for creating a starter document that companies can use to make decisions about how to design AI-enabled products and services that interact with children.

Panellists:

Nick Allen, Ann Swindells Professor of Clinical Psychology, University of Oregon

Christopher Payne, Senior Manager, Government and Public Affairs EMEA, LEGO Group

Michael Preston, Executive Director, Joan Ganz Cooney Center at Sesame Workshop

Yalda Uhls, Executive Director, UCLA Center for Scholars and Storytellers, University of California, Los Angeles

Moderator:

Chloe Poynton, Principal and Co-Founder, Article One

Soliciting multistakeholder support

The reality of corporate governance differs depending on the type of business in question. Businesses whose primary consumer are children face a distinct set of opportunities and challenges relative to general market providers that also interact with them.

Businesses need external stakeholders' support to understand how to design products that contribute positively to child development in a demonstrable way. Providing a detailed roadmap of steps that *enhance development through design* would be helpful.

Some ideas for multistakeholder support include consulting experts, testing products directly with children, collaborating with the state on progressive policies and working within the industry to create professional standards and corresponding mechanisms of accountability.

Incorporating developmental science

Product design could also be enhanced by considering aspects of human development, such as: 1) developmental tasks that occur at each stage of life; and 2) steps that assess whether technology is helping or harming these tasks. One possibility is to align key development processes (i.e. "fundamental learning tasks") with product use, ultimately striving to evaluate whether a product is beneficial for children given the science. This would require designing channels to translate efforts from the academic to business context for operability.

Strengthening corporate responsibility

The tension between choosing the most responsible business practice and optimizing for commercial viability is constant. The goal is to strengthen corporate responsibility, including by building momentum behind certain key performance indicators and developing structures that enable responsibility. Businesses have primarily focused on responsibility in terms of safety while, in the context of digital technology, shifting the focus to development could create additional value. This idea supports a theme that emerged throughout the workshop on changing the AI perception from one of fear to one of opportunity.

Another issue to consider is whether using a "carrot" instead of a "stick" can motivate changes in corporate behaviour. These changes have typically been prompted by public shaming or remedial measures designed to correct for reputational damage. Ultimately, the stick approach of influencing corporate behaviour is not sustainable because it is reactive rather than proactive.

Breakout group discussions

Internal processes

Key points:

Establishing and aligning across teams. Companies need to evaluate if teams should be added or restructured to support decision-making with respect to children and AI. They should also establish deliberate processes to align coordinated action between the different divisions of the company.

- The value of establishing a Research Department should be considered, focused broadly on issues of product, technology and society to the company ecosystem. Even in companies that have already established this team, more could be done to integrate cutting-edge findings into engineering systematically.
- The importance of ensuring that AI technologies are placed in the right hands means the issue of how to position salespeople is critical. For example, should dedicated training sessions for entire sales teams be organized on evaluating targeted buyers, or should management include a dedicated, ethics-focused decision-maker?

Instituting an ethics review board. A company ethics review board could be beneficial if the incentives are correctly aligned. One suggestion is to establish an "ethics code" to support the board's work.

- The ethics code would align with other policy-making processes within a company and act as a guiding and interpretative force.

- The board might adapt the code from a supply-chain model to allocate responsibility in situations where companies are using AI acquired from or selling to other providers.
- If the notion of “ethics” is too vague to apply concretely, examining the moral principles of human rights could be helpful, as they have been defined at the international level and include a detailed set of responsibilities and mechanisms for evaluating compliance.
- The idea of human rights does not translate as readily as ethics in business, but there might be value in working through this complexity to elevate its position as a guiding concept for business.
- One challenge with peer-to-peer messaging is to ensure the information delivery seems authentic; children are good at spotting messages that feel contrived relative to the mode of content delivery.

Promoting positive messages. The media is and always will be prone to reporting exaggerated stories that stoke fear. The focus needs to shift on promoting the positive potential of AI and accurate research.

Taking children into consideration to create a development-focused design code. The corporate governance charter should specifically include the needs of children with respect to AI, as distinct from AI principles more generally. The charter should:

- Incorporate feedback from stakeholders speaking strictly in the interest of the child and who have no interests in the business
- Include the results of analyses of the potential harmful effect that algorithmic delivery or other products can have on child audiences
- Take into consideration vulnerabilities based on the developmental science and stipulate the need to monitor the impact of technology use to respond to the manifestation of harmful effects that may emerge over time. These elements would anchor a *development-focused design code*.

– Internet access is key to providing learning resources and other growth opportunities for more marginalized children. Any measures to limit a child’s access to certain platforms should be evaluated alongside the value of opportunities lost.

– The focus in this area is on identifying how to make the most impact within an imperfect system. This may at times be a question of education, and at others of persuasion.

– Another element is understanding how to use media systems; for example, establishing journalism fellowships in certain domains could help organically expand the universe of issues being reported on regarding AI’s impact on children.

Educating parents. Initiatives to educate parents on the present-day, scientific realities of AI are needed. Parents typically know less about digital technologies than their children, which means their ability to exercise meaningful control is challenged.

– Groups might stage “upward interventions” in which teens educate their parents on technology platforms and use. This would help to achieve both a healthy respect for adolescent autonomy and parental ability to exercise meaningful behavioural control.

– Children see digital technology primarily as an opportunity, while parents tend to see it as harmful. Using the term “generation gap” to frame the misunderstanding of the impact of technology could help parents see things in a different way.

– Efforts to provide reliable, expert information should focus on reaching parents along all race, class and socio-economic lines, to empower children more equitably.

Public education

Key points:

Reaching young people. The group considered how to reach young people with positive messages regarding the potential influence and applicability of AI in their lives.

- Peer information is highly valuable to children and especially adolescents, who are often resistant to information delivered from authority figures. Young people are much less engaged with typical media, including movies, than with new media, such as communication hosted by peer-to-peer ecosystems like YouTube and Instagram.
- One option is to work with a consortium of youth influencers to help them disseminate positive messages about AI. By way of analogy, the Norman Lear Center at the University of Southern California works to inject messages into Hollywood film-making that promote positive behaviour.

Consumer protection

Key points:

The group focused on AI as it applies to media construction and consumption online. A regulatory framework for consumer protection was judged by certain participants to be too outdated to apply to the online space, where children not only consume but also create and market content. As advertisers, young people market themselves as the product, on the one hand, and endorse consumer products, on the other.

Considering creators, consumers and advertisers. The definition and corresponding treatment of children need to be reframed to include their commercial roles in a world of monetized user-generated content.

- Taking children seriously as content producers who fuel the media economy supports new solutions that balance protection with agency and peer-to-peer communication.
- Current notions of consumer protection assume that only adults have agency and that, therefore, any protective or moderating measure must be imposed by adults. Yet children are in the best position to help structure the rules regarding their engagement in multifaceted roles online.
- Devising policy that focuses on promoting a child's healthy engagement with the digitized world rather than on protection would help companies to acknowledge the realities of technology use and to support the empowerment of young people on their platforms as both creators and consumers.

Defining the limits of protectionism. A protectionist mindset may unduly restrict access to knowledge. What needs protection and the justification for providing it must be clear.

- A regulatory framework would be beneficial in that companies would understand the parameters they are operating within.
- The specific product and industry in question must be kept in mind when considering whether regulation is need or other means should be used, such as working through media and education to foster awareness and change public perceptions.

Assessment and evaluation

Key points:

Aligning design with development. In the focus on aligning product design with child development, is it possible to determine whether the AI is working to advance or limit progress?

- Processes are needed to document behaviour, emotions and patterns of sociability to provide meaningful data for measures of impact.
- With regard to data collection to evaluate impact, companies must also think about mitigating privacy concerns. Practices should be in place to ensure data anonymization while allowing the aggregation of insight relevant for measuring impact. These practices should be clearly articulated and conveyed to customers to promote corporate trust and establish sustainable practices for data integrity.

Accounting for marginalized groups. When measuring “developmental impact,” deliberate care should be taken to ensure that its effects on vulnerable and marginalized groups of children are taken into account.

- Depending on their data sources and evaluation methods, companies run the risk of drawing inferences regarding the impact that miss considering children at risk adequately. Related to their data practices, companies should think proactively about how bias might be embedded into their machine learning systems.
- Companies should think about setting up processes to identify and measure the impact of the technology in question on vulnerable populations as distinct from the general group.

Consolidating multiple models of impact. A single model of developmental impact is too limiting; the science is evolving and scientists have yet to understand the many aspects of development. At the same time, it would be counterproductive for each company or industry to define its own concepts of development.

- To strike a balance, one goal is to produce a charter that demonstrates some flexibility in defining developmental parameters, alongside clear priorities that reflect consensus as appropriate.
- This approach also supports innovation on the part of companies, causing the greater likelihood of corporate buy-in, since overly broad requirements can restrict iteration and growth.
- The potential of companies to experiment with new product concepts that might further learning and development objectives must not be hindered.

Upholding auditing and reporting. Irrespective of the impact model used, an audit and reporting structure must be in place for any corporate charter to function at full capacity.

- A company's size and resources are among the factors that determine its practical auditing capabilities.
- The goal of audits and reporting is not necessarily to penalize, but to provide an accurate assessment that will enable the company to make targeted improvements.
- Setting parameters externally helps to create a single set of industry-wide standards focused on children and AI. A single model of impact would facilitate entrenching best practices across industry and creating norms that govern and are useful by indicating market expectations.

Public policy

Plenary overview: Panel presentations

The discussion focused on the lack of current government policies regarding AI and the opportunities to move forward. The objective was to think about policy guidelines, including how countries might structure and implement them to shape sovereign laws on a global scale.

Panellists:

Alexa Koenig, Executive Director, Human Rights Center, Berkeley Law, University of California, Berkeley

Angela Vigil, Partner and Director, Pro Bono Practice, Baker McKenzie

Steve Vosloo, Policy Specialist, Digital Connectivity, United Nations Children's Fund (UNICEF)

Examining the current policy landscape

Published in May 2019, Generation AI key contributor law firm Baker McKenzie conducted a multi-jurisdictional survey evaluating AI in existing law, and particularly as it might apply to children and toys. The primary finding was that no laws currently exist that directly regulate or mention AI. This positions the present as a golden moment to set the tone for policy in this field. Starting from a blank slate, stakeholders are poised to create a positive culture around AI and law, *if they act now*.

“

There are no laws regulating AI anywhere.

”

The survey also found that due to the current dearth of relevant policy, the judiciary is eager for laws that directly apply to AI. Current decisions are being made based on pre-existing local laws in relevant domains, such as product and privacy laws and guidance from international treaties. This is problematic because the defining characteristics of AI were not contemplated when existing laws were enacted, making it likely they are ill-suited to the regulatory task.

The survey recommends adopting sector-specific laws for AI rather than broad regulations that might fail to capture variance and important balancing concerns in AI applications. The youth voice should also be sourced and its vantage and creativity used to inform forward-thinking democratic laws that reflect the unique experience of those they are intended to govern.

Ensuring fairness in machine learning

Algorithms that exercise control represent relationships that are not always easy to explain or understand. As advances in deep learning processes are leveraged to create more complex algorithmic models that strive for greater accuracy, interpretability becomes even more challenging.

Another challenge of machine learning is to ensure that the model is complex enough to render the decision that it is being asked to make. If training data does not accurately represent the entire group of people on which an algorithm operates, the decisions that it renders vis-à-vis individuals at the margin can be biased and inaccurate. Overcoming bias when AI interacts to render decisions that impact children is even more important due to the potential for negative developmental effects that could harm a child for life.

Shaping AI educational policy

The use of AI in the classroom is set to double over the next few years, which provides stakeholders with an opportunity to move ahead by shaping policy today. In addition, children are learning through media platforms online, outside of public educational settings. Robust AI educational policy will consider learning from different entry points to address the diverse aspects of how society is shaping and teaching today's young minds.

The issues pertaining to the classroom are myriad. AI is being brought into teaching tools in a way that alters how students learn; the potential for these applications should be evaluated and directed towards enhancing learning outcomes through a developmental frame. The opportunities and risks of personalized learning tools should also be evaluated to ensure maximum benefit, particularly regarding their potential to nurture growth for children with special needs or who are otherwise disadvantaged relative to their peers. One of the most significant outcomes of shifting the societal narrative to focus on positive AI could be its potential to equalize educational opportunities.

Another issue that bears consideration is surveillance in schools. As learning becomes increasingly digital and schools continue to gather more data on youths' movements and learning activities, structures that protect youngsters against the potentially harmful or prejudicial misuse of their data should be put in place. While there may also be opportunities for helpful interventions, for example if an AI could detect that a child is at risk of self-harm, the risk of adverse effects is also high. How schools should proceed in these sensitive contexts deserves special policy consideration.

Breakout group discussions

Laws and regulation protecting children

Key points:

Scrutinizing existing laws. The consensus from the multi-jurisdictional survey conducted by Baker McKenzie was that the regulatory landscape is scant.

- Insofar as applicable, existing laws focus only on preventing harm and enabling parental rights.

- None focuses on leveraging technology to promote developmental objectives, which presents a key opportunity and mandate for policy going forward.

Ensuring informed public participation. The type of information that allows the public to make informed decisions on AI policies with respect to children is important. Priorities include:

- Understanding data use and its benefits
- Demystifying AI and changing the narrative from being anchored in stories of fear and harm
- Providing simple analogies that translate complex concepts about AI
- Focusing on specificity with respect to potential harms to promote tailored decision-making.

Considering data ownership. Data in the context of children and AI-enabled ecosystems raises certain questions including:

- What ownership do people have over their data and how might data ownership apply differently to individuals as private citizens versus public figures?
- Personal data encompasses everything from fluffy social media posts to immutable biometric data; how should privacy be applied differentially with respect to data as both an external and internal asset?
- Under what circumstances can data ownership rights be waived?

Government protection

Key points:

Ensuring equitable data practices. One's existence as a child is no longer ephemeral: children are producing content and data about their movements in real time that is liable to become a permanent part of their digital footprint going forward.

- Under the GDPR, “the right to be forgotten” aims to protect individuals in the development of their autonomy by preventing them from being stigmatized for actions performed in the past. Under its mandate, internet search providers must remove certain information about individuals if requested.
- The practical implementation of the right to be forgotten could provide an unfair advantage to children with resources to cultivate online reputation management, thereby enhancing conditions for inequality.

- The tracking of children's disciplinary measures could have a negative impact on employment prospects and college admissions. In view of the implications, can policy be enacted to limit how this type of data is collected and used?

Restricting content exposure. Limiting content exposure should be considered relative to the different types of content and the corresponding assessment of harms. Despite what the policies state or intend, the ability to restrict content exposure is limited once a child has internet access.

- Currently, advertisers use many workarounds to reach children, related to the ways they engage online. For example, YouTube Kids has rules against targeted advertising, but the reality is that most children bypass age disclaimers to establish a presence on YouTube's main platform. The issue of how to limit targeted advertising to children is even more salient.
- Less discussed but equally if not more relevant is the issue of exposing youngsters to interpersonal content. What are the negative implications of children constantly consuming photos and stories of the lives of their peers and public figures? Pending such analysis, is regulation necessary or should companies handle these issues directly?

Mandating internet access. Should access to the internet be protected as a fundamental right of childhood?

- In many educational settings, children are expected to have internet access to conduct research and complete assignments. The internet also provides educational opportunities and resources where none are provided by parents or institutions.
- How would mandating internet access confront the desire of certain parents to raise their children in offline environments?
- The questions remain regarding when to defer to parents as proxies for the best interests of the child and when the state should establish a baseline for social participation.

Defining the “best interests of the child”. Entrusting parents with the task of brokering a child's opportunities in the digital context can be problematic when, in many instances, parents understand little about the technology, often even less than their children.

- Children are native to digital technologies in a way their parents will never be. This highlights the importance of bringing children's voices into the legal definition of the “best interests of the child” in the area of AI.

- Relying entirely on children to set the terms is also problematic and impractical. This highlights the need to incorporate alternate sources of authority (for example, those supported by developmental psychology) to create policies that determine how children will interact with and use AI.
- The former social contract between parents, society and children dictated that parents controlled education in the home and the state controlled education at school. Learning is now happening online, where neither parents nor schools are intermediaries.
- Can a new intermediary be established that controls the validity and quality of knowledge and is aligned with public interest? This may bridge the gap between parents and children in serving the “best interests of the child” as it applies to new forms of learning online.
- Constantly interacting with government allies and working to build and develop these relationships over time
- Determining a clear request when communicating with legislators
- Leading with incentives for legislation proposed
- Leading with the public consultation process to promote trust and transparency
- Selecting a single jurisdiction for testing to establish best practices and iterate based on lessons learned
- Leveraging targeted external expertise to inform the development of specific issues
- Creating a proposal outlining incentives for various groups in government (who may not ordinarily work together) to collaborate on the initiatives proposed, thereby also furthering their own objectives
- Incorporating equalizing factors that address vulnerable groups of children, such as more severe penalties for disproportionate effects on these groups.

Science to policy

Key points:

Enabling common understanding. It is important to include explainers into initial discussions with stakeholders when drafting policy, who can translate scientific concepts into policy, to best identify and lead with points promoting multistakeholder governance.

- The concept of a common vocabulary is important to understanding and drafting policy.
- Policy should be subjected to a multistakeholder review process.
- Some thought should go into deciding at what stage children and other relevant stakeholders not directly involved should be included in the process.

Recognizing developmental science in policy. The methods for recognizing developmental stages in policy are varied. The goal is to arrive at policy that maps the main learning stages to help companies create metrics that optimize for well-being.

- Currently, policies that apply variably to different groups of children do so by age.
- The objective of incorporating developmental stages into policy is to encourage product development that optimizes the needs of children at each phase.

Structuring policy guidelines. Policy guidelines could incorporate model legislation parameters that offer creative ways to influence legislation. They might also recommend:

Policy guidance for AI and child rights

The discussions in this section were framed in collaboration with UNICEF and examined how ethics and rights-based concepts should be incorporated into policy guidelines for government-level policy makers.

Key points:

Considering ethics and human rights. Ethical principles are a useful starting point for policy conversations but often lack specific content and accountability mechanisms related to human rights.

- The concept of ethics as an operating principle is more straightforward and has therefore taken greater hold in the AI and related business communities.
- What are the advantages of working within an ethics vs human rights framework, given the relative merits and drawbacks of each approach?
- The human rights framework offers a base from which to categorize priorities, including, for example, the right to education, the freedom of thought and the “maximum development of the child” as noted in the UN Convention on the Rights of the Child, to which developmental science can be incorporated to delineate specific priorities.

Ensuring algorithmic accountability. Optimizing for well-being is a central tenet of policy guidelines on how to regulate business. When considering accountability and rights, policy might also focus on the need to:

- Conduct an early-impact assessment of the algorithms that will be used in public decision-making processes, and prescribing steps to mitigate potential harms
- Establish transparency measures that document how algorithms are designed and used, to create accountability in decision-making.

Creating context. Human stories are essential to provide useful guidance and context for policy-making.

- Policy guidance can benefit from case studies included to illustrate factual applications of general principles.
- Consideration should be given to adapting policy guidelines to the cultural context while establishing a minimum baseline of actions needed to protect children as individuals on a universal scale. It may be appropriate to limit the role of culture where established interests of the child are concerned.

Outcomes

This workshop aimed to survey the broad range of issues relevant to two key deliverables: 1) a set of public policy guidelines geared towards government-level decision-makers focused on children and AI; and 2) a corporate governance charter that guides companies developing AI as they design products and services with child audiences in mind. The next stages of work will seek to draw from the group's contributions and synthesize their conclusions to draft the two documents.

The workshop was ambitious in scope; this report serves to capture and explore the multiple dimensions of the conversation in crude fashion. It is important to note that not every point is relevant to every social instance or AI application. This extensive review of the discussion points seeks to provide interested stakeholders with a reference point to advance on both dimensions of this work. The key points related to the workshop themes, and the structure and priorities for each deliverable, are summarized below:

Workshop themes

- The narrative on AI should shift from a focus on fear and potential harm to an emphasis on opportunities and possibility.
- Where privacy is concerned, decision-makers should conduct a balanced evaluation of countervailing interests to determine the priorities for accomplishing the objectives at stake, including finding new ways of enabling privacy that recognize the importance of data to the advancement of AI.
- Parents may not always be in the best position to determine the relationships between their children, AI and technology use. This dynamic frames a threefold need to educate parents, identify new intermediaries and find new ways to bring children as stakeholders into the policy-making process.
- Developmental science should be incorporated as a baseline for both company and country guidelines, while recognizing that using one static model of impact will unduly limit the evolving science and product innovation.
- Ethics may be predominant in the discussion on governing AI, but principles of international human rights law offer negotiated content and standards that an appeal to ethics lacks. The relative strengths of each approach should be considered in framing actionable recommendations.

Corporate governance charter

- The charter should reflect a set of positive commitments and include a roadmap to promote knowledge and understanding among corporate decision-makers and to structure operational processes throughout the echelons of the company to support the outcomes prescribed by the document.
- Steps should be taken with industry leaders to set up and finance an auditing entity or administrative body that works to enforce commitments.
- The focus of the charter is to enhance development through design.

Policy guidelines

- Little can be effectively generalized from the application of existing laws to regulate AI; new theories that account for its defining characteristics are needed to construct appropriate policy measures.
- A key dimension of creating policy at the state level is determining ways to educate citizens, thereby enabling meaningful democratic participation with respect to legislative policy-making concerning AI.
- Since the AI policy landscape is currently a blank slate, stakeholders have a major opportunity to shape the narrative on the development of AI if they act now.

Workshop participants

Carla Aerts, Director, Tmrw Digital, United Kingdom

Nick Allen, Ann Swindells Professor of Clinical Psychology, University of Oregon, USA

Maria Luciana Axente, Programme Driver, Artificial Intelligence, Technology & Investments & AI for G, PwC Strategy&, United Kingdom

Brent Barron, Director, Public Policy, Canadian Institute for Advanced Research, Canada

Danielle Benecke, Senior Associate, IPTech, Baker McKenzie, USA

Caroline Burnett, Professional Support Lawyer, Baker McKenzie, USA

Kay Firth-Butterfield, Portfolio Head, Artificial Intelligence and Machine Learning, World Economic Forum

Ronald Dahl, Director, Institute of Human Development, University of California, Berkeley, USA

Rebecca Finlay, Vice-President, Engagement and Public Policy, Canadian Institute for Advanced Research, Canada

Anna Goldenberg, Chair in Biomedical Information and AI, SickKids, Canada

Mizuko Ito, Director, Connected Learning Lab, University of California, Irvine, USA

Erica Kochi, Co-Founder, UNICEF Innovation, United Nations Children's Fund (UNICEF), New York

Alexa Koenig, Executive Director, Human Rights Center, Berkeley Law, University of California, Berkeley, USA

Emma Morris, Director, Global Policy, Family Online Safety Institute (FOSI), USA

Kelsey Noonan, Program Strategist, Pivotal Ventures, USA

Candice Odgers, Fellow, Child and Brain Development, Canadian Institute for Advanced Research, Canada

Christopher Payne, Senior Manager, Government and Public Affairs, EMEA, LEGO Group, Denmark

Chloe Poynton, Principal and Co-Founder, Article One, USA

Michael Preston, Executive Director, Joan Ganz Cooney Center at Sesame Workshop, USA

Javier Raygada, Software Engineer, Facebook, USA

Neeta Saran, Senior Attorney, Microsoft Corp., USA

Stephanie Seale, Consultant, Partnerships and Outreach, Office of the High Commissioner for Human Rights (OHCHR), Geneva

Jordan Shapiro, Fellow, Brookings Institution, USA

Karen Silverman, Partner, Latham & Watkins, USA

Yalda Uhls, Executive Director, UCLA Center for Scholars and Storytellers, University of California, Los Angeles (UCLA), USA

Angela Vigil, Partner and Director of Pro Bono Practice, Baker McKenzie, USA

Steve Vosloo, Policy Specialist in Digital Connectivity, United Nations Children's Fund (UNICEF), USA

Robert Wicks, Software Engineer, Facebook, USA

World Economic Forum contact

Eddan Katz

Project Lead, Artificial Intelligence and Machine Learning

Eddan.katz@weforum.org

Aaina Agarwal

Project Specialist, Artificial Intelligence and Machine Learning

Aaina.agarwal@weforum.org



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744

contact@weforum.org
www.weforum.org