COALITION FOR INNOVATION

AI Blueprint for the Future

Coalition for Innovation, supported by LG NOVA

Jami Diaz, Director Ecosystem Community & Startup Experience
William Barkis, Head of Grand Challenges & Ecosystem Development
Sokwoo Rhee, Executive Vice President, LG Electronics, Head, LG NOVA

Coalition for Innovation Co-Chairs

Alex Fang, CleanTech Chair
Sarah Ennis, AI Chair
Alfred Poor, HealthTech Chair

Authors

Adrien Abecassis, Johnny Aguirre, John Barton, Ann M. Marcus, Olivier Bacs, Taylor Black, Micah Boster, Mathilde Cerioli, Carolyn Eagen, Sarah Ennis, Annie Hanlon, Christina Lee Storm, Andrew Yongwoo Lim, Jess Loren, Refael Shamir, Svetlana Stotskaya

Senior Editor, Alfred Poor
Editor, Jade Newton

October 2025

Preamble

The Coalition for Innovation is an initiative hosted by LG NOVA that creates the opportunity for innovators, entrepreneurs, and business leaders across sectors to come together to collaborate on important topics in technology to drive impact. The end goal: together we can leverage our collective knowledge to advance important work that drives positive impact in our communities and the world. The simple vision is that we can be stronger together and increase our individual and collective impact on the world through collaboration.

This "Blueprint for the Future" document (henceforth: "Blueprint") defines a vision for the future through which technology innovation can improve the lives of people, their communities, and the planet. The goal is to lay out a vision and potentially provide the framework to start taking action in the areas of interest for the members of the Coalition. The chapters in this Blueprint are intended to be a "Big Tent" in which many diverse perspectives and interests and different approaches to impact can come together. Hence, the structure of the Blueprint is intended to be as inclusive as possible in which different chapters of the Blueprint focus on different topic areas, written by different authors with individual perspectives that may be less widely supported by the group.

Participation in the Coalition at large and authorship of the overall Blueprint document does not imply endorsement of the ideas of any specific chapter but rather acknowledges a contribution to the discussion and general engagement in the Coalition process that led to the publication of this Blueprint.

All contributors will be listed as "Authors" of the Blueprint in alphabetical order. The Co-Chairs for each Coalition will be listed as "Editors" also in alphabetical order. Authorship will include each individual author's name along with optional title and optional organization at the author's discretion.

Each chapter will list only the subset of participants that meaningfully contributed to that chapter. Authorship for chapters will be in rank order based on contribution: the first author(s) will have contributed the most, second author(s) second most, and so on. Equal contributions at each level will be listed as "Co-Authors"; if two or more authors contributed the most and contributed equally, they will be noted with an asterisk as "Co-First Authors". If two authors contributed second-most and equally, they will be listed as "Co-Second Authors" and so on.

The Blueprint document itself, as the work of the group, is licensed under the Creative Commons Attribution 4.0 (aka "BY") International License: https://creativecommons.org/licenses/by/4.0/. Because of our commitment to openness, you are free to share and adapt the Blueprint with attribution (as more fully described in the CC BY 4.0 license).

The Coalition is intended to be a community-driven activity and where possible governance will be by majority vote of each domain group. Specifically, each Coalition will decide which topics are included as chapters by majority vote of the group. The approach is intended to be inclusive so we will ask that topics be included unless they are considered by the majority to be significantly out of scope.

We intend for the document to reach a broad, international audience, including:

- People involved in the three technology domains: CleanTech, AI, and HealthTech
- Researchers from academic and private institutions
- Investors
- Students
- Policy creators at the corporate level and all levels of government

Chapter 7:
What AI Owes Children: A New Blueprint for User-Centered Beneficial Innovation

Authors: Mathilde Cerioli, Adrien Abécassis

Imagine what social media might look like today if, in 2008, we had asked child development experts some basic questions. Should we expose young girls to constant appearance-based filters during identity formation? Should emotionally charged or violent content be algorithmically reinforced for boys during critical windows of social learning? Should children and adults interact freely on the same platforms—with no meaningful supervision or safeguards? And what if, instead, we had designed for long-term wellbeing: promoting empathy, critical thinking, healthy connection, and mechanisms that prevent addiction?

We didn't ask then, but with generative AI reshaping digital experiences once again, we have another chance at designing a tech environment that prioritizes children's developmental needs and fundamental rights. The iRaise Alliance's mission is to do exactly this: build the frameworks, standards, and collaborations needed to design AI with children's development, rights and futures at the center from the very beginning.

Overview

The iRAISE Alliance (International Research-driven Alliance for AI Serving Every child) is a global, multi-stakeholder initiative launched in 2025 to fundamentally shift how AI systems are designed, implemented, and governed for children. Grounded in child development, neurosciences and child rights, the Coalition brings together governments, researchers, tech companies, NGOs, and civil society to build an ecosystem that proactively supports children's well-being in digital environments.

This white paper outlines the potential AI yields for young children, the developmental risks, as well as the Alliance's unique approach to bridging the systemic gaps in industry, research, and regulation. By connecting research, design, policy, and public awareness in an integrated model, this initiative aims to redefine beneficial AI from the ground up—placing children at the center of design, not at the margins of risk management.

We invite partners across sectors to join this growing movement—contributing expertise, investment, and support to shape a future where AI protects and empowers the next generation.

Context

During the first 25 years of life, the human brain undergoes rapid and profound changes that shape each individual's cognitive and socio-emotional capacities. As a result, the experiences children and adolescents are exposed to play a decisive role in determining who they become and what they are capable of achieving. This developmental window makes them especially receptive to opportunities—but also particularly vulnerable to external influences, including those mediated by digital technologies.

By dramatically altering children's environments through increasingly ubiquitous digital interfaces, AI raises an essential question: Are we ensuring that this new environment supports their growth rather than disrupts it?

AI fundamentally reshapes our world and offers significant new opportunities for expression, connection, and learning. **For children and**

adolescents, it could unlock more equitable access to education globally and across socioeconomic divides, while providing personalized learning experiences tailored to each child's unique needs and abilities. AI can also help children realize their rights—especially in contexts where those rights are under threat—such as their rights to education, freedom of expression, or access to information, culture, and participation in decision-making.

Researchers and child advocacy organizations are raising two major concerns when it comes to child development: cognitive and emotional overreliance on AI. When children engage with AI-supported learning tools, the technology promises enhanced educational access and personalized support. At the same time, however, it risks undermining their capacity for independent and critical thinking. The line between both is fine, and only careful design—informed by seasoned experts in learning, education, and cognitive development—can ensure AI becomes a force for good.

Closely related is a second concern: the rise of parasocial relationships between children and AI—one-sided emotional attachments to media figures, fictional characters, or, increasingly, artificial intelligence. Unlike traditional media, conversational AI responds and directly engages with children, often using anthropomorphic, emotionally charged designs that simulate empathy. This makes such agents especially powerful—and potentially harmful—for developing brains. Poorly designed AI risks distorting children's understanding of social relationships, weakening emotional resilience, and interfering with neurological reward systems. These risks intensify in high-exposure contexts — such as among children experiencing isolation, trauma, or inconsistent caregiving — where AI chatbots may begin to substitute for real human connection.

By developing products from the outset with children in mind, we can move beyond risk mitigation to creating systems that genuinely serve and strengthen their development.

Current limitations

However, the current system, left to its own devices, will fail to meet this imperative; just as social media has failed younger generations by not adapting to their developmental needs or by not placing their wellbeing first.

At a structural level, these risks are compounded by systemic limitations. First, **there is a deep disconnect between research and product development**. AI product teams must make constant decisions—how to build, adapt, and integrate models for children and adolescents—often without access to the developmental expertise required to do so responsibly. Meanwhile, researchers—bound by methodological caution—are hesitant to issue concrete recommendations when data is limited. This creates a vacuum: those most qualified to provide guidance often won't, and those making critical decisions may lack the developmental literacy to do so well.

Another key issue is the difference in operational speed and ethical standards between research and industry. **Scientists must adhere to strict ethical protocols and cannot test potentially harmful scenarios, while companies face no such requirement before releasing child-facing products at scale.** This disparity in standards and timelines exacerbates the mismatch: research proceeds deliberately, while industry moves at the pace of the market.

On a global scale, this ethical divide also impacts regulation. **Most regulatory frameworks are reactive, intervening only after harm is shown.** And because research cannot ethically study harms in advance, a paradox emerges: we must prove that children are already being harmed before regulators can act. This "wait and see" approach slows progress and redirects resources away from innovation and toward damage control. Instead of asking how AI can best support development, efforts are wasted trying to document harm post hoc—squandering time and delaying urgently needed course correction.

These dynamics have created an ecosystem where child development is treated as collateral—not as

a constraint or design principle. Even well-intentioned innovations, aimed at increasing access or engagement, can backfire: optimizing for short-term gains at the cost of deeper learning, or addressing loneliness through systems that cultivate emotional dependency on non-human agents.

The true cost is opportunity lost. **As long as we treat child safety as a regulatory afterthought, we forfeit the chance to build AI tools that are not only safe—but profoundly developmental, equitable, and transformative.**

Our unique approach

Within the iRAISE alliance, we are redefining how AI products are designed, through a simple yet radical premise: AI should be designed with and for children from the outset. This approach requires collaboration from the beginning; only by responsibly developing products from the start can we achieve more with and through AI, focusing our resources on creating beneficial products rather than spending them on harm mitigation.

When products are truly developed to serve their users—and include a broad range of stakeholders from the very beginning, from those most directly impacted, such as children, to subject-matter experts, technologists, and systems-level policy thinkers—the need for correction, regulation, and punitive measures is significantly reduced.

To address those limitations, Everyone.AI and the Paris Peace Forum launched the Beneficial AI for Children Coalition in February 2025 at the Paris AI Action Summit, following a year of global consultation and a pilot convening in San Francisco. This work is now framed as the **iRAISE alliance, defined as a multi-actor initiative bringing together governments, academic researchers, technology companies, NGOs, and civil society actors.** The Alliance brings together representatives from over a dozen governments (Bulgaria, Chile, Costa Rica, Denmark, France, Iceland, Luxembourg, Mexico, Norway, Senegal, Togo, Uruguay); leading AI companies such as Google, OpenAI, Anthropic, and Hugging Face; and international organizations, such as the United Nations and UNESCO. It also includes over 20 NGOs and civil society organizations, like Common Sense Media, 5Rights Foundation, Joan-Ganz Cooney Center, The Alan Turing Institute; as well as renowned researchers (Stuart Russel, Isabelle Hau, Michael Preston, David Harris, Florence GSell, Sonia Livingston), and leading research labs from top academic institutions, including Stanford Social Media Lab, Access to Knowledge for Development Center (American University in Cairo), Social Brain Science (Zurich), Boston Children's Digital Wellness Lab, Connected Learning Lab (Irvine). This diverse, cross-sectoral participation ensures the Alliance reflects a truly global perspective—firmly grounded in both policy and practice. We are especially intentional in involving organizations with experience in amplifying children's voices. Co-design is not optional—it is foundational. Children must be participants in shaping the tools designed for them, not just recipients of their outcomes.

Our approach is built to correct the very disconnects identified above. While the full architecture is still in early stages of implementation, the model we are building connects research, design, policy, and public engagement in a cohesive, mutually reinforcing ecosystem. Rather than treating these domains as separate, we are developing a framework where they inform and strengthen one another from the outset. Cross-sector collaboration is central to this vision. **We are establishing spaces—such as confidential, multi-stakeholder labs—that bring together product developers, child development experts, and researchers to exchange findings, challenge assumptions, and define age-appropriate priorities for safer, more developmentally-aligned AI.** These labs are intended to generate practical design guidance while also surfacing research questions that can guide future academic inquiry.

A transdisciplinary research hub is being developed in parallel to explore how AI affects children's cognitive and socio-emotional development. By engaging experts from neuroscience, psychology, linguistics, sociology, and computer science, this work will ensure that AI design is informed by the latest evidence on child development. **These insights will ultimately translate into age-specific design**

standards that can guide both product teams and policy discussions.

At the same time, we are initiating dialogue with policymakers, governments, and international organizations to ensure that regulation evolves alongside innovation. The aim is to create shared frameworks that anticipate and shape the development of child-centered AI, rather than reacting to harm after it has occurred.

Public awareness and knowledge-sharing will be an integral part of our broader approach. We are building a digital knowledge platform to make tools, findings, and real-world practices accessible to researchers, companies, and decision-makers. Convenings and workshops will support learning, co-creation, and the emergence of a shared language across disciplines and sectors.

Current efforts

Our first milestone was the publication of a foundational research report "The Future of Child Development in the AI Era" published in 2024 and developed through close consultation with both child development and AI experts. The work is now recognized as a valuable contribution to the field and continues to serve as the basis for our ongoing efforts. It exemplifies the rigorous and collaborative approach we believe is necessary to responsibly shape AI's role in children's lives.

This ecosystem is still under construction—but the foundation is already proving strong. **In December 2024, we hosted our first closed-door workshop, gathering over 50 thought leaders from across the ecosystem**. This event demonstrated not only a collective willingness to collaborate, but also the value of creating space for honest, cross-disciplinary exchange. Participants gained a clearer understanding of each other's realities, constraints, and motivations—laying the groundwork for more integrated, coordinated solutions.

The formal launch of the coalition at the AI Action Summit in February 2025 built on this momentum. Despite being established in just a few months, the Alliance rapidly brought together some of the most influential actors in the space—evidence of both the urgency and shared commitment to building an AI future that respects and uplifts the next generation.

The first Child-AI Lab is scheduled to begin in the fall of 2025 and is expected to yield an early draft of emerging best practices. Since the launch, several leading research labs from prestigious universities have reached out to join the initiative, underscoring its potential to become a hub for international collaboration. In parallel, additional governments are preparing to join the Alliance ahead of the Paris Peace Forum in October 2025, where key updates and further announcements will be made.

This growing engagement across sectors shows that momentum is building—and that a shared, proactive, and child-first approach to AI is both necessary and possible.

Call to Action

Children are already engaging with AI systems every day, yet most are still not designed with their developmental needs in mind—a critical gap with long-term consequences. The iRAISE Alliance is building the foundation for a new approach: one where child development, rights, and agency are embedded into AI systems from the start. To realize this vision, we need the combined influence, expertise, and support of actors across sectors.

If you are driving investment decisions, your backing can accelerate the development of scalable, research-informed solutions—and position you at the forefront of responsible innovation. If you shape policy, your engagement can help align emerging regulation with child-centered standards that are globally coherent and locally effective. If you build products, this is your opportunity to lead by example—setting new benchmarks for trust, safety, and long-term value by designing for and with children in mind. If you conduct research, your knowledge can guide real-world impact—bridging the gap between science and design while shaping how AI serves child development across contexts. And if you are a

parent, teacher, designer, or engaged citizen, your awareness, advocacy, and day-to-day choices all help build demand for systems that prioritize children's well-being.

This work is underway, but its success depends on collective ownership and sustained support—including the resources to move from vision to implementation. The systems we shape today will influence generations to come. Join us in ensuring that AI grows up with children—not ahead of them.

How to Join and Apply This Work

The iRAISE Alliance is designed as a collaborative platform where governments, companies, researchers, NGOs, and individuals can take tangible steps to embed child-centered design into AI development and governance. Engagement can begin at any scale and deepen over time.

- **Start where you are**: As parents, caregivers, teachers, or mentors, consider how AI tools—and the way children use them—interact with each child's literacy and interests. Choosing age-appropriate systems and encouraging critical use through ongoing conversations can measurably support healthy development.
- **Assess your impact**: For organizations and institutions, review how your AI tools or services influence children, whether directly or indirectly. This includes products not explicitly marketed to children but still used by them.
- **Include child development expertise**: Integrate developmental science into the earliest design phases. This can mean inviting advisors to prototype reviews, collaborating with educators, or working with organizations experienced in amplifying children's voices.
- **Pilot child-centered design**: Test age-appropriate AI in high-impact areas such as education, content moderation, and digital literacy. Measure not only usability but also its impact on attention, motivation, empathy, and resilience.
- **Join the Alliance's collaborative work**: Contribute to open working groups or share research through the Global Knowledge Platform to align with emerging, research-informed design standards.

By engaging in these actions, you help move beyond harm prevention toward unlocking AI's full potential to support every child.

Author (In order of contribution)

Dr. Mathilde Cerioli, Chief Scientist, everyone.ai
Dr. Mathilde Cerioli is the Chief Scientist and cofounder of everyone.ai, a nonprofit dedicated to anticipating and educating on the opportunities and risks of AI for children. She holds a Ph.D. in Cognitive Neuroscience and a master's degree in Psychology, with a research focus on how AI intersects with cognitive and socioemotional development in children, adolescents, and young adults. In May 2024, she published the influential report Child Development in the AI Era, examining the potential impact of emerging technologies on cognitive and socioemotional development.

Adrien Abecassis, Executive Director for Policy at Paris Peace Forum
Adrien Abecassis is a French career diplomat and a former senior advisor to the President of France

(2012–2017). He has held academic fellowships at Harvard University and UCLA and is currently serving as Chief Policy Officer of the Paris Peace Forum.

For more information about the Coalition for Innovation, including how you can get involved, please visit coalitionforinnovation.com.

View the Next Chapter