

This forum features practitioner perspectives on designing technologies for and with communities. We highlight compelling projects and provocative points of view that speak to both community technology practice and the interaction design field as a whole. — **Sheena Erete, Editor**

Child-Centered AI:

Contextualizing Principles and Design in HCI

Ge Wang, University of Illinois Urbana-Champaign, **Ayça Atabey**, University of Edinburgh, **Kaiwen Sun**, Indiana University Bloomington, **Grace C. Lin**, Massachusetts Institute of Technology, **Samantha-Kaye Johnston**, University of Melbourne, **Kruakae Pothong**, 5Rights and London School of Economics and Political Science, **Cara Wilson**, University of Edinburgh, **Lachlan D. Urquhart**, University of Edinburgh, **Jason C. Yip**, University of Washington, **Jun Zhao**, University of Oxford

AI technologies are increasingly integrated into the daily lives of children, encompassing tools that include learning platforms, interactive robots, and generative AI systems. This has spurred growing interest in HCI and related fields that emphasize the development of AI systems that prioritize children's best interests. In the summer of 2023 and 2024, we gathered a group of 67 expert researchers across HCI and related disciplines at CHI workshops. By leveraging their expertise, we offer a multidisciplinary perspective on child-centered AI.

While the HCI community has a long history of designing for children, a disconnect persists between the design community, policy guidelines, and the actual designs implemented in practice. Emerging policies—UNICEF's policy guidance on AI for children, general comment No. 25 on children's rights in relation to the digital environment from the United Nations Convention on the Rights of the Child, and the Council of Europe's *Handbook for Policy Makers on the Rights of the Child in the Digital Environment*—focus more on articulating the big picture of what technologies should do for children. Regulations, such as the EU AI Act and the UK Online Safety Act, emphasize child protection from harm resulting from the operation and use of digital technologies, including AI. These policies and regulations have yet to speak the language of designers and developers who are working on concrete

and practical actions that they can take to account for children's unique developmental needs and rights in AI system design.

In the realms of policies and regulations, the notion of child-centered AI is often overlooked or oversimplified. In contrast, HCI experts work closely with children, using a child-centered approach and offering valuable insights into how AI can be effectively designed to serve children's best interests, addressing their unique needs and vulnerabilities. This expertise presents a crucial, yet often underexplored, perspective on AI integration, going beyond what policy documents typically capture.

RUNNING THE WORKSHOP

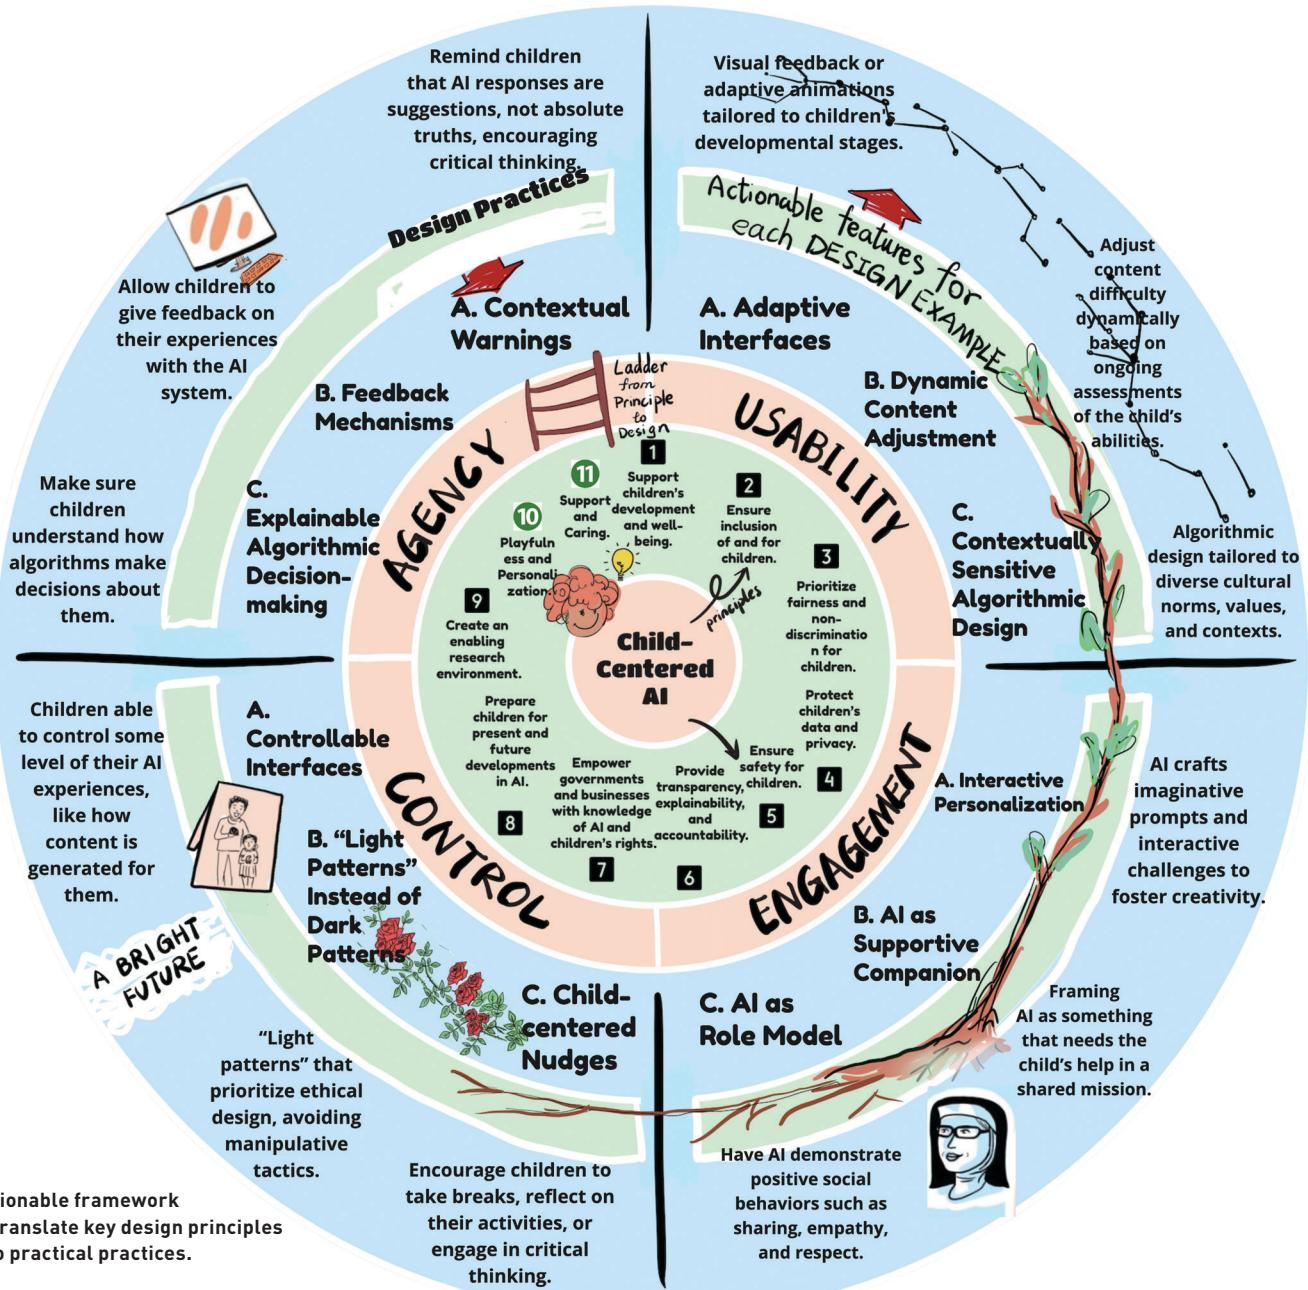
We conducted two full-day hybrid workshops at CHI '23 and CHI '24. We chose a two-year workshop approach to allow for iterative refinement of our methods and to broaden the range of participant perspectives over time. We followed an identical methodological approach both years.

Insights

- Child-centered AI lacks focus on long-term impacts and robust evaluation metrics.
- Expertise remains siloed across domains, hindering the development of multidisciplinary approaches to child-centered AI.
- To bridge these gaps, we propose an actionable design framework for child-centered AI, informed by insights from our workshop discussions.

Participants were recruited through various channels, including our website, HCI-related e-lists, social media, and word of mouth. Potential applicants were required to submit a position paper detailing their professional background, experiences with child and AI technologies, and a discussion on child-related AI technologies they have worked on or are familiar with in industry or research. We aimed for a diverse pool, prioritizing international and institutional diversity, and a balance of expertise across different facets of the research topic. For both years, we emailed all registered participants beforehand to obtain their consent for using their workshop outputs for research purposes. Based on these criteria, we selected 67 participants from 13 countries and regions.

The full-day workshops began with participants presenting their past research and visions on child-related AI during five-minute talks. Beyond academic presentations, the workshops incorporated collaborative mechanisms, including discussions via Discord and structured group activities facilitated by co-organizers. For seamless interaction, participants were divided into in-person and online groups, with all activities documented on a shared Miro board. The workshops featured four activities: defining child-centered AI using keywords, establishing guiding principles through brainstorming, assessing current AI practices by identifying positive and negative examples, and codesigning AI solutions for children by refining existing AI applications based on prior discussions. Groups remained consistent throughout



the workshops to build on previous insights, and their outcomes were presented to all participants.

KEY PRINCIPLES FOR CHILD-CENTERED AI

The most fundamental principle emphasized by our participants was the importance of supporting children's development and growth in child-centered AI systems, aligning closely with existing research in child-computer interaction. However, as noted in previous studies [1], the focus on children's developmental needs is often limited to learning and entertainment

systems. While participants provided examples of good design, including tailoring support or adapting to children's specific needs in areas such as literacy development, there is still a lack of attention to the long-term impacts of child-centered AI systems and an absence of robust evaluation metrics to assess these effects. This limits our ability to measure how these systems influence children's developmental trajectories over time.

Meanwhile, participants highlighted the overlooked importance of playfulness and care as central elements in AI design. Although ethics and

fairness dominate current guidelines, playfulness and care are critical for fostering engagement and supporting positive developmental experiences. The absence of these considerations underscores a broader disconnect between theoretical principles and practical implementations in AI design.

Another critical gap lies in the limited focus on children with special needs. Current frameworks, such as UNICEF's guidelines, fail to address the unique challenges this underrepresented group faces, resulting in a lack of inclusive systems that accommodate diverse developmental requirements.

Expanding these frameworks to explicitly incorporate considerations for children with special needs is vital for equitable design outcomes.

Finally, participants noted that expertise in designing for children remains siloed across cognitive development, learning sciences, ethics, and other domains. This fragmentation hampers the creation of cohesive, multidisciplinary approaches to child-centered AI. Breaking down these silos and fostering collaboration among researchers and practitioners can enable more nuanced and tailored design practices that better address the diverse needs of children.

Apart from the design principles themselves, our discussions provided key insights into ways to integrate ethical design principles more deeply into the child-centered design community. For instance, participants emphasized the challenge of balancing ethical considerations, such as privacy, with the need for adaptive designs. Designing AI systems that adapt to a child's developmental stage often requires collecting sensitive data, including behavioral patterns and demographic information, raising significant ethical and privacy concerns. Even within the design principles, tensions emerge, particularly when conflicting priorities come into play. For example, while adaptive systems can enhance children's learning and emotional development, they also carry the risk of fostering overattachment if not carefully managed. Prior research has proposed strategies to navigate these conflicts. One such example is designing systems that prioritize real-world social interactions while incorporating safeguards to limit excessive use [2]. Other studies emphasize the importance of providing children with choices and fostering their autonomy to help balance personalization with privacy and developmental goals [3]. Thoughtfully designed AI systems that integrate these approaches can help mitigate risks while ensuring an appropriate balance between conflicting design priorities. Finally, while there is a long tradition of making sure that ethical practices are used when children are involved in research, less attention has been given to embedding ethical decisions

directly into the design of technologies. In particular, developers feel there is little consumer recognition of the value of such meticulous design work, which reduces incentives for incorporating these features. These challenges suggest the need for more robust frameworks and practices that make ethical considerations a central part of the design process.

BRIDGING POLICY AND DESIGN: TOWARD AN ACTIONABLE FRAMEWORK ON CHILD-CENTERED AI

Discussions in our workshops point toward an actionable framework to translate key design principles into practical practices.

Usability. This aspect of child-centered AI design focuses on enabling children to effectively and efficiently use AI. Accessible, intuitive, and age-appropriate interfaces should align with HCI research on usability for young users, ensuring that AI-generated content matches the child's developmental stage and navigation prevents cognitive overload. We encourage future designers to draw directly from the insights shared by our participants, which highlight several effective practices: adaptive interfaces, such as visual feedback or animations tailored to developmental stages; dynamic content adjustment, where AI adapts the difficulty level based on assessments of a child's abilities; and contextually sensitive algorithmic design, which tailors algorithms to reflect diverse cultural norms, values, and contexts.

Engagement. This theme explores how AI systems can stimulate children's creativity, interaction, and social behaviors. This aligns with research on social learning and peer-like interactions, which demonstrates that children thrive through collaborative and creative processes. Effective design in this area could involve interactive personalization,

where AI fosters creativity through imaginative prompts, or framing AI as a supportive companion that collaborates with children in shared missions. Another promising approach is designing AI to embody positive behaviors (e.g., empathy, sharing, respect) and serve as a role model.

Control. This theme focuses on empowering children to actively control AI. Future design components could include approaches such as controllable interfaces that enable children to adjust aspects of their AI experience—for example, influencing how content is generated—while incorporating light patterns and ethical design principles to avoid manipulative tactics. Similarly, nudges can be used to remind children to take breaks, reflect on their activities, and critically evaluate how AI influences their digital experiences. These practices align with findings from participatory design and AI personalization research, which emphasize the importance of giving users control over features such as content recommendations and privacy settings. This could foster a sense of ownership and shared authority, allowing children to actively shape their interactions with AI systems.

Agency. The theme centers on allowing children to move from simply controlling AI to critically reflecting on its decisions and functions. This stage invites children to engage with AI thoughtfully, understand its processes, and question its outcomes. To achieve this, effective design practices could include contextual warnings—for example, reminders that AI responses are suggestions rather than absolute truths—encouraging children to critically assess the information presented. Feedback mechanisms can also play a role, enabling children to provide input on their experiences with AI systems, fostering a sense of ownership and collaborative improvement. Additionally, transparent algorithmic decision making, where systems explain how decisions are made, can ensure that children understand the logic and biases underlying algorithmic outcomes, supporting informed interactions. These practices are backed by research on transparency and bias prevention in AI, which stresses the importance of

HCI experts can help create systems that protect children while fostering their autonomy.

presenting multiple perspectives and acknowledging errors to foster critical thinking and trust [4]. Moreover, studies have shown that involving caregivers in the process can strengthen trust building and encourage reflective dialogue, helping children better understand the ethical and societal implications of AI systems [5].

CONCLUSION

The four themes we outlined—usability, engagement, control, and agency—align closely with an intended progression of children’s engagement with AI. The process begins with basic interactions, where children learn to use AI effectively, and advances to deeper levels of control and critical reflection, where they actively shape and evaluate their interactions with AI. This developmental trajectory aligns with children’s cognitive and social growth, illustrating how AI design can support their evolving capabilities and autonomy. It also resonates with prior work on algorithmic autonomy, which emphasizes the importance of fostering critical thinking, independent decision making, and the ability to reflect on and challenge AI-driven outcomes. We encourage future researchers and designers to use this framework as a bridge between ethical design principles and practical design practices. HCI practitioners can integrate ethical considerations into the design process from the outset, ensuring that AI systems are not only effective but also respectful of children’s rights and privacy. By developing frameworks that combine ethical guidelines (e.g., transparency, privacy, and data minimization) with child-centered usability, HCI experts can help create systems that protect children while fostering their autonomy.

Our work suggests that child-centered AI has the potential to establish its own distinct road map, complementing the traditional focus of child-computer interaction on usability, accessibility, and interaction design. While child-centered AI is still in its early stages, this article introduces new challenges and complexities unique to AI for children, providing a foundation

for future exploration and development. Our work consolidates a research agenda, but further investigation and collaboration will be essential to address the evolving intricacies of designing AI for children, ensuring that future solutions effectively support their growth, development, and well-being.

ENDNOTES

1. Wang, G., Zhao, J., Van Kleek, M., and Shadbolt, N. Informing age-appropriate AI: Examining principles and practices of AI for children. *Proc. of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, 2023, Article 536, 1–29; <https://doi.org/10.1145/3491102.3502057>
2. Wang, G., Zhao, J., Van Kleek, M., and Shadbolt, N. ‘Treat me as your friend, not a number in your database’: Co-designing with children to cope with datafication online. *Proc. of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, 2023, Article 95, 1–21; <https://doi.org/10.1145/3544548.3580933>
3. Badillo-Urquiola, K., Smriti, D., McNally, B., Golub, E., Bonsignore, E., and Wisniewski, P.J. Stranger danger! Social media app features co-designed with children to keep them safe online. *Proc. of the 18th ACM International Conference on Interaction Design and Children*. ACM, 2019, 394–406; <https://doi.org/10.1145/3311927.3323133>
4. Dworkin, G. *The Theory and Practice of Autonomy*. Cambridge University Press, 1988.
5. Friedman, M.A. Autonomy and the split-level self. *The Southern Journal of Philosophy* 24, 1 (1986), 19–35.

¶ **Ge Wang** is an assistant professor of computer science at the University of Illinois Urbana-Champaign. Her research focuses on the intersection of HCI, human-centered AI, and usable security and privacy.
→ wangge@illinois.edu

¶ **Ayça Atabay** is an interdisciplinary researcher and tutor at the University of Edinburgh and at the Digital Futures for Children center at the London School of Economics and Political Science. She works on AI regulation, data protection and AI use in education, fairness of data processing, and age-appropriate design of edtech with and for children.
→ aycaatabay@ed.ac.uk

¶ **Kaiwen Sun** is an assistant professor at Indiana University Bloomington. Her research spans HCI, usable security and privacy, and

human-centered AI, focusing on populations such as children, families, and older adults by designing and developing safe, transparent, inclusive, and privacy-protective solutions.

→ kaiwsun@iu.edu

¶ **Grace C. Lin** is a learning scientist and assessment designer at the Massachusetts Institute of Technology, dedicated to enhancing children’s and students’ capacity to learn. Recently, her work has focused on child-centered AI, AI education for children, the integration and implications of AI technology on human skills, and its influence on human relationships.
→ gcl@mit.edu

¶ **Samantha-Kaye Johnston** is a McKenzie Research Fellow at the University of Melbourne, researching critical thinking, autonomy, and AI-enabled edtech. She leads projects on AI in education, co-edits Oxford University Press’s AI in Society series, and is a visiting researcher at the University of Oxford’s Department of Computer Science.
→ samanthakaye.johnston@unimelb.edu.au

¶ **Kruakae Pothong** is a researcher at 5Rights and visiting research fellow in the Department of Media and Communications at the London School of Economics and Political Science. Her current research focuses on child-centered design for digital services and children’s education data.
→ kruakae@5rightsfoundation.com

¶ **Cara Wilson** is a lecturer in children and technology at the University of Edinburgh.
→ cara.wilson@ed.ac.uk

¶ **Lachlan D. Urquhart** is a senior lecturer in technology law and HCI at the Edinburgh Law School at the University of Edinburgh. He is founder and director of the Regulation and Design Lab. His main research interests are in the sociotechnical aspects of designing, living with, and regulating emerging information technologies.
→ lachlan.urquhart@ed.ac.uk

¶ **Jason C. Yip** is an associate professor focusing on digital youth at the Information School and an adjunct assistant professor in the Department of Human-Centered Design and Engineering at the University of Washington. He builds innovative technologies for new collaborations and examines how current technological trends already influence family collaborations around learning.
→ jcyip@uw.edu

¶ **Jun Zhao** is an Oxford Martin Fellow and founder of the Oxford Child-Centred AI Design Lab in the Department of Computer Science at the University of Oxford. Her research focuses on investigating the impact of algorithm-driven decision making on everyday lives, with a particular emphasis on families and young children.
→ jun.zhao@dcs.ox.ac.uk

