The second workshop on Child-Centred AI design (CCAI)

The Second Workshop on Child-Centred AI Design (CCAI)

Ayça Atabey
ayca.atabey@ed.ac.uk
School of Law, Centre for Data,
Culture & Society. University of
Edinburgh
UK

Ge Wang[*]
ge.wang@cs.ox.ac.uk
Department of Computer Science,
University of Oxford
UK

Samantha-Kaye Johnston
samantha-
kaye.johnston@cs.ox.ac.uk
Department of Computer Science.
University of Oxford
UK

Grace C. Lin
gcl@mit.edu
Scheller Teacher Education Program |
The Education Arcade. The
Massachusetts Institute of Technology
USA

Cara Wilson
cara.wilson@ed.ac.uk
Centre for Research in Digital
Education. University of Edinburgh
UK

Lachlan Urquhart
lachlan.urquhart@ed.ac.uk
Edinburgh Law School. University of
Edinburgh
UK

Jun Zhao
jun.zhao@cs.ox.ac.uk
Department of Computer Science.
University of Oxford
UK

ABSTRACT

AI-powered solutions are increasingly woven into the fabric of children's digital worlds. They're found in interactive toys, home automation systems, everyday apps, and various online services that young users engage with. As we look ahead, it's almost certain that the prevalence of AI in tools and platforms designed for kids will grow, given AI's ability to offer rich, tailored, and dynamic experiences. However, the nuances of how these AI-centric platforms cater to children and how they can be optimized to meet the unique needs of younger users remain largely underexplored. Building on the momentum from our inaugural CCAI workshop at CHI 2023, our aspirations for this year's workshop include: (1) deepening the discourse on the essence of AI that prioritizes children, (2) focusing on actionable strategies to operationalise child-centered AI design principles into practice, and (3) cultivating an ever-growing collective of professionals passionate about the future of child-focused AI innovations.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**.

KEYWORDS

child-computer interaction, Child-Centred AI design

[*]Corresponding author

$_H owtooperationalisechild-centredAIinpractices$"

1 BACKGROUND

AI technologies are becoming more integrated into children's daily routines. This encompasses assistive learning platforms, interactive robots, intelligent home devices, and generative AI/LLMs platforms like ChatGPT that children engage with regularly [1, 25]. Such AI tools offer advantages such as customized educational guidance from intelligent tutoring systems [9, 21], and online safety measures to protect children [14, 24]. In the realm of education, AI provides personalized content suggestions [10], while social robots animate characters in relatable, human-like ways [4, 13], some of which are particularly beneficial for vulnerable groups of children [3, 30]. Nevertheless, despite the claimed benefits AI systems could have in children's lives, their use also raises significant concerns given the many long term uncertainties with their implementation [6]. For example, AI systems could create potential biases against certain groups [5, 12], such as vulnerable groups of children who might face risks and unfairly treated in a way that impacts their wellbeing and daily lives [1, 25]. Similarly, inscrutability and unpredictability could inadvertently cause children to be exposed to harm in content filtering systems in ways that were difficult to anticipate or predict, such as those crafted by malicious adversaries [1]. Moreover, children are among those at greatest risk of privacy-related harms when Internet-connected smart home devices designed without children's unique needs in mind collect children's data that could

have long-term impact on them [2, 23]. Parents also expressed concerns regarding children's safety using smart home technologies when AI systems lack safety guardrails for child users, such as lacking granular access control [23, 26]. Going forward, AI systems are likely to become even more ubiquitous in children's lives with the rise of generative AI/LLMs; as such, there is a pressing need to address these developments through a CCAI lens and explore how to translate CCAI into the design of AI systems for children.

Last year, we hosted the first CCAI (child-centered AI) workshop at CHI 2023. The event garnered 100+ submissions, drawing in 40+ participants both on-site and virtually. The results of the workshop successfully culminated in the 'Hamburg Declaration on Child-Centered AI', which was signed by all participants and identified critical issues surrounding the design for child-centered AI. This declaration was presented at the House of Lords and also inspired a special issue call on Child-Centered AI at IJHCS. Our focus for the first CCAI workshop was on understanding, conceptualising, and framing child-centered AI, engaging in discussions about its meaning and foundational design principles. This year, our focus will shift to the practical implementation of child-centered AI. With the increasing integration of new AI technologies into children's daily lives—such as the introduction of generative AI in classrooms and LLMs in personal robots for kids—it is imperative to re-examine the meaning of child-centered AI in these contexts. And we believe there is a pressing need to transform high-level design principles into actionable guidelines and best practices for practitioners and developers.

We believe CHI is an ideal venue for this workshop. In fact, we have already seen great interest in topics around child-centred AI in CHI. For instance, varying work on designing for and with children appear each year in different subcommittees such as *Learning, Education, and Families* [11, 15, 28], *Design* [16, 18, 27], *Games and Play* [7, 17, 19], and *Privacy and Security* [8, 22, 29]. We recognize a compelling need to unite individuals from various domains within CHI who share an interest in designing AI from a child-centered perspective, along with designers and developers at the forefront of creating products and services for children. Expanding on our prior endeavors to foster interdisciplinary conversations—by engaging experts from fields like education, child development, and public policy, ensuring a multifaceted exploration of factors, methodologies, and viewpoints that influence AI and human interactions—we aim this year to actively involve industry professionals, educators, and hands-on practitioners who work directly with children, enriching the discourse even further. By doing so, we hope to create a holistic dialogue that not only highlights theoretical insights but also emphasizes real-world applications, bridging the gap between research and on-the-ground practices in child-centered AI.

2 GOALS OF THE WORKSHOP

This workshop is designed to navigate the intricacies of translating child-centered AI design principles into real-world applications. Existing policies, like the Online Safety Bill [20], introduce substantial operational challenges when we attempt to embed these principles into the actual design and rollout of child-centric AI systems. These challenges are further amplified by the evolving landscape of legal and policy frameworks. In spite of the progression in these frameworks, there is a notable gap in identifying pragmatic solutions to these challenges. As we explore operationalization, we'll touch on aspects ranging from conceptual to methodological to technical to ethical. This encompasses frameworks, adaptable evaluation methods, and concrete design considerations. To guide our exploration, we have outlined a series of questions, which, while not exhaustive, serve as thought starters:

- What specific guidelines, standards, or benchmarks are necessary for AI developers and companies to adhere to child-centered policies? Additionally, how might market regulators and certification bodies weave child-centred AI criteria into their evaluations and approvals?
- How can we establish effective feedback loops involving AI practitioners, children, caregivers, policymakers, educators, and students to consistently refine child-centered AI policies, system designs, and classroom tools?
- How can design teams efficiently involve child psychologists, educators, assessment specialists, and children in AI design and testing? And, in terms of AI user interfaces (UI) and user experiences (UX) for children, what best practices can ensure safety, fairness, and transparency?
- Given the potential for higher costs or extended timelines, how can we motivate app developers and product teams to prioritize child-centered AI design considerations?
- What measures can educational institutions take to ensure that AI tools used in classrooms adhere to child-centered designs? Additionally, how should curricula evolve to comprehensively educate students on the utilization, advantages, and potential risks of AI?
- What are the implications, both positive and negative, of embedding child-centered AI into popular products and platforms, and how can we effectively convey these to parents and guardians?

3 ORGANISERS

We believe our organizing committee is well-suited to conduct this workshop, given both the diversity of our disciplines as well as a common interest grounded in child-centred AI design. Our team comes from a variety of backgrounds in the field of HCI, User Experience (UX), Human-Robot-Interaction, Education, Law, as well as Regulations and Policy.

Ayça Atabey will serve as the lead co-organiser for this workshop. She is PhD student at Edinburgh University. She received PhD Enrichment Award from Alan Turing Institute, where she currently works on Child-Centred AI. She is an affiliate at the Centre for Data, Culture & Society of the Edinburgh Futures Institute. She conducts interdisciplinary research at the intersection of IT Law and Human-Computer Interaction. She also works as a Consultant on AI and Data Protection at UN Women and a researcher at IT Law Institute. She has also worked at the Digital Futures Commission at 5Rights Foundation focusing on EdTech regulation and child rights by design.

Ge Wang is a PhD student at the University of Oxford, Department of Computer Science. Her research focuses on Child-Computer/AI interaction, and investigates the algorithmic impact on families and children, exploring the potential for designing more age-appropriate and autonomy-supportive AI for families. She led the first CCAI workshop @ CHI 2023, and has experience co-organising several workshops, seminars and forums, and served as a volunteer in several previous conferences and workshops.

Samantha-Kaye Johnston is a Researcher at the University of Oxford, Department of Computer Science and a research affiliate at the Berkman Klein Centre for Internet and Society at the Harvard Law School. Her research is focused on how people read, think and create in digital environments. She is currently leading the Digital Autonomy Machine experiment and Project Amplify. Samantha is a mixed-methods researcher with a background in conducting large scale initiatives across schools and community organisations in different countries to support global conversations around people's skill development as they navigate the digital world.

Cara Wilson is a Lecturer in Children and Technology and a member of the Children and Technology Group at the Institute for Education, Community and Society. With a background in Human-Computer Interaction and Interaction Design, she explores how to better support inclusion and diversity in the development of technologies which pervade children's lives, for example, working with neurodiverse children in primary school contexts. Through creative, explorative and participatory design methods, her research endeavours to support the strengths, voice and agency of individuals who are often overlooked in the process of design.

Grace C Lin is a Learning Scientist and Assessment Designer at MIT with over 10 years of experience in research projects focusing on building students' capacity to learn. Her past and present projects include a practitioner-oriented early childhood measures repository, ed tech games and apps, and project-based learning curricula in math and AI that incorporate playful assessments.

Lachlan Urquhart is a Senior Lecturer in Technology Law and Human-Computer Interaction. He is Director of the Regulation and Design Lab and Co-Director of the Scottish Research Centre for Intellectual Property and Technology Law (SCRIPT) and Centre for Research on Information, Surveillance and Privacy. His main research interests are in the socio-technical aspects of designing, living with, and regulating emerging information technologies. He has been awarded £7m in external research projects and published 60 papers across law, ethics and computing.

Jun Zhao is a Senior Research Fellow at the University of Oxford, Department of Computer Science. Her research focuses on investigating the impact of algorithm-based decision makings upon our everyday life, especially for families and young children. For this, she takes a human-centric approach, focusing on understanding real users' needs in order to design technologies that can make a real impact. Currently, she is leading the KOALA project and the ReEnTrust project. She work closely with schools, children, families

as well as technologists for children, to understand the technological, societal and regulatory challenges that we are facing, to inform national and international policymakers, technology designers and families. Jun is also the UK regional leader for the 100 Brilliant Women in AI and Ethics global initiative to promote diversity and equality in this critical research area.

4　WORKSHOP LOGISTICS

Below we describe the logistics for our workshop, this includes our **pre-workshop plans**, the **mode of workshop**, the **website, Discord channel and asynchronous engagement**, as well as our **accessibility statement**.

Pre-Workshop Plans. Our pre-workshop plans serve three goals: **advertising** (to raise awareness and receive strong submissions), **building community**, and **recruiting extra expert reviewers**. **First**, the call for participation will be distributed via HCI-, UX-, AI, Edu- and Policy- for children-related mailing lists. We will further utilise the mailing list we had from our previous year's CCAI workshop (which attracted 100+ submissions). In addition to this, our workshop from the previous year successfully evolved into a special issue call on Child-Centered AI at IJHCS=. We will also invite selected publications to present at our workshop. **Second**, in terms of community building, apart from the organisers' personal academic networks, we also plan to utilize the engagement through advertisements on social media to expand our community. **Third**, drawing from last year's data, we received over 100 submissions spanning a variety of disciplines within HCI, including areas such as Education and Law. Our review process will adopt a journal-style approach. Initially, the committee members will conduct a preliminary review to filter out submissions unrelated to the workshop topics. Subsequent to this, each submission will be evaluated by two externally recruited reviewers. Our priority is to ensure a diverse range of perspectives and representation, aiming to make the workshop as inclusive and fair as possible.

Workshop Mode: Hybrid. To promote equitable participation, we will host a hybrid workshop. We arrived at this decision after consulting with various CHI stakeholders, taking into account visa restrictions and sustainability considerations. The hybrid format enhances global participation, as it diminishes the relevance of travel costs and visa concerns for remote attendees. Additionally, we hope it would promote sustainability for Hawaii. Meanwhile, we aim to support both the experience of our in-person and remote participants. Remote participants will engage in interaction with in-person participants through a variety of means including our website, Discord channel, Miro boards as well as Zoom (which has been proven to be quite effective and engaging from our previous year's experience), as we outline below.

Website, Discord Channel, Zoom, Miro and Asynchronous Engagement. The workshop website can be found at https://www.ccai2023.org/. While it currently displays content from last year's workshop, we will update it upon acceptance of this proposal. It

contains the call for participation including dates and author instructions. We will provide background information on the topics, actual links to interesting news on HCI and CCAI, as well as, the background of each organizer. Our website will further provide a rich source of information and engagement for the workshop to enable the access of keynotes, expert panel discussions, paper presentations, downloadable proceedings, as well as portals to our group design activities. Given the hybrid nature of the proposed workshop, Discord will host our participants virtually. We will also have a hybrid registration desk to ensure registered participants get access to workshop related activities. Taken together, the website, Discord channel, Miro boards afford effective asynchronous engagement as these tools proved to support discussions and save records of the presentations and design activities based on our prior experiences. Beyond asynchronous avenues, we will use Zoom for live presentations and Q&A sessions.

Accessibility Statement. We are committed to providing an inclusive environment and we will do our best to accommodate requests for any special assistance from our participants. All materials will be provided in an accessible format (e.g., subtitles and transcriptions will be provided throughout the workshop). Authors submitting to our workshop will be strongly encouraged to work on improving the accessibility of their papers, including adding figure descriptions, in compliance with the SIGCHI's Guide to an Accessible Submission (https://sigchi.org/conferences/author-resources/accessibility-guide/).

5 WORKSHOP STRUCTURE

This will be a one-day workshop consisting of approx. 40 participants (both in-person and online) and the organizers. The workshop will be conducted in a hybrid format, combining traditional paper presentations and ideation and prototyping activities. Approximately **2 weeks before the workshop**, we will send out links to all participants to introduce themselves in the Discord channel and to have access to workshop materials. Those intending to attend in-person are invited to prepare a three-minute presentation and forward us their slides. Meanwhile, online participants are asked to submit a three-minute video recording of their work. The deadline for both slides and video submissions is three days prior to the workshop.

The workshop (see Table 1) will be devoted to a half-day keynote speech and paper presentation session, and a half-day group activity session. The keynote speech and all paper presentations will happen both in-person and in Zoom while the discussion happens on dedicated channels in Discord. Based on previous experience, this combination not only promoted a smooth experience (without cluttering the chat on Zoom calls) but also allowed for asynchronous engagement. Moreover, speakers appreciated being able to continue the conversation threads in Discord even after their talks are over.

In the **morning sessions** of the workshop, we will begin with a brief *introductory session* that aligns participants with the workshop goals, outlines key activities, and introduces the organizers. Next, we will feature a *keynote speech* from a distinguished speaker who is at the forefront of child-centred AI and HCI. Last year, we were honored to host Baroness Beeban Kidron, who discussed the policy implications of child-centered AI. This year, our aim is to invite speakers from the industry to provide insight into the practical aspects of child-centred AI. The rest of the morning will include sessions for *paper presentations.*

The **afternoon sessions** of the workshop mainly involves group activities including group discussion and group design activities. It kicks off with an *introductory session* that aligns participants with the goal of the group activities, and the material to use for group discussion and design. Then, *group discussion* takes place, in which the participants are invited to explore critical issues on translating child-centered AI principles into practices, contributing towards a deeper understanding of what future child-centred AI design might mean for children. Discussion topics will be crowd-sourced and curated by the organising committee, but mainly around topics such as *What are the design principles around child-centered AI?*; *How do we operationalise these principles into practices?*. We will gather topics through surveys before the workshop. These topics can be shared before the workshop to allow for group formation before we go live. This way, we can minimize coordination challenges, which will help us keep the workshop on time. Breakout rooms on Zoom will be provided (max 6 people per room). After the group discussion activity and coffee break, participants will be invited to the *group design* activity, participants will form into groups of 5-6, and explore the design space of child-centred AI by envisioning and creating low-fidelity paper prototypes, operationalising the previously discussed design principles into practice. To promote the interaction between in-person and remote participants, the design activities will be mainly conducted through Miro (https://miro.com/), which is a tool for online board collaboration. We will make sure each group contains an approximately equal number of in-person and remote participants.

6 POST-WORKSHOP PLANS

We have a three-part plan regarding the post-workshop outcomes. The first is to do with producing achievable results and contents. Our workshop is committed to recognizing and disseminating the valuable contributions of its participants. We intend to compile and publish a collection of the submitted papers as workshop proceedings. To ensure accessibility and widespread distribution, we are considering reputable platforms such as https://ceur-ws.org and ArXive, utilizing their report numbers. This approach not only elevates the visibility of participants' work but also adds to the collective knowledge in the field. We also plan to invite authors of standout submissions from the workshop to further develop and submit their papers to a special issue in a selected reputable journal, mirroring our approach from last year, when we expand our workshop into a special issue call on Child-Centered AI at IJHCS.

Secondly, we plan to synthesise and analyse the results gathered from our group discussion and group design activities into a report on the workshop outcome first, which may be transformed into publishable results to selected journals or magazines. Last year, our collective efforts culminated in a joint report, the 'Hamburg Declaration on Child-Centered AI', which was signed by all participants and presented at the House of Lords.

Time	Session
2 weeks before the workshop	Participants introduce themselves in the Discord Channel, and have access to workshop materials.
9:00 - 9:20	Introduction of workshop organisers, participants, topics, and goals
9:20 - 10:20	Keynote Speech by invited speaker and Q&A session
10:20 - 10:30	Coffee break
10:30 - 12:00	Paper Session
12:00 - 13:30	Lunch Break
13:30 - 14:00	Introduction of group activities, topics, and goals
14:00 - 15:20	Group Activity 1: Group Discussion
15:20 - 15:30	Coffee Break
15:30 - 17:00	Group Activity 2: Group Design
17:00 - 17:30	Wrap-up
17:30	Drinks & Dinner (optional)

Table 1: Tentative workshop schedule (time is in local timezone)

Thirdly, a central goal of this workshop is community building for researchers and practitioners in this area. So, we plan to continue the conversation on Discord with our participants. We also plan to use the website as an archival repository of workshop contents and new resources to foster continuous conversations and attract new community members. Last, if there is a critical mass of interested participants, we will explore opportunities to transform the workshop/forum into a new conference in the future (similar to how FAT* workshops lead to ACM FAccT conference).

7 CALL FOR PARTICIPATION

AI systems and algorithms have become central to children's digital landscapes, seamlessly integrating into connected toys, smart home IoT devices, and daily-use apps. Their prevalence is attributed to their ability to enhance user experiences, making them adaptive, compelling, and personalized. During our first CCAI workshop at CHI 2023, we focused on conceptualizing child-centered AI and highlighted its significance. This year, our second CCAI workshop will pivot to the hands-on application of these conceptual foundations in child-centered AI. We will investigate how next-gen technologies like generative AI/LLMs and personal robots can be designed from a child-centric perspective. We invite papers, up to 4 pages excluding references, covering: predominant challenges, methodologies employed, foundational principles, practical safeguards, and insights for translating child-centered AI concepts into practice. Submissions should conform to the CHI Extended Abstract format and be submitted via our workshop's portal (https://www.ccai2024.org/). The submissions will be reviewed by the workshop organizers and program committee. All accepted papers will be published through our website and presented at the workshop, provided at least one author attends the workshop and registers at least one day of the conference.

REFERENCES

[1] 2019. Generation AI: Establishing Global Standards for Children and AI. http://www3.weforum.org/docs/WEF_Generation_AI_%20May_2019_Workshop_Report.pdf

[2] 2020. Age appropriate design code. https://ico.org.uk/media/for-organisations/guide-to-data-protection/key-data-protection-themes/age-appropriate-design-a-code-of-practice-for-online-services-2-1.pdf

[3] Jonnathan Berrezueta-Guzman, Vladimir Espartaco Robles-Bykbaev, Iván Pau, Fernando Pesántez-Avilés, and María-Luisa Martín-Ruiz. 2021. Robotic technologies in ADHD care: Literature review. *IEEE Access* 10 (2021), 608–625.

[4] Ginevra Castellano, Iolanda Leite, André Pereira, Carlos Martinho, Ana Paiva, and Peter W. Mcowan. 2014. Context-Sensitive Affect Recognition for a Robotic Game Companion. *ACM Trans. Interact. Intell. Syst.* 4, 2, Article 10 (June 2014), 25 pages. https://doi.org/10.1145/2622615

[5] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*. 797–806.

[6] Emma Day, Kruakae Pothong, Ayça Atabey, and Sonia Livingstone. 2022. Who controls children's education data? A socio-legal analysis of the UK governance regimes for schools and EdTech. *Learning, Media and Technology* (2022), 1–15.

[7] Thomas Dylan, Gavin Wood, Abigail C. Durrant, John Vines, Pablo E. Torres, Philip I. N. Ulrich, Mutlu Cukurova, Amanda Carr, Sena Çerçi, and Shaun Lawson. 2020. Designing IoT Resources to Support Outdoor Play for Children. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376302

[8] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. "Money Makes the World Go around": Identifying Barriers to Better Privacy in Children's Apps From Developers' Perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 46, 15 pages. https://doi.org/10.1145/3411764.3445599

[9] Ahmed Fadhil and Adolfo Villafiorita. 2017. An Adaptive Learning with Gamification and Conversational UIs: The Rise of CiboPoliBot. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization* (Bratislava, Slovakia) *(UMAP '17)*. Association for Computing Machinery, New York, NY, USA, 408–412. https://doi.org/10.1145/3099023.3099112

[10] Mohamed Heni Frikha, Tarek Zlitni, and Nadia Bouassida. 2020. Towards a Recommendation System for Children's Animated Movies Based on Chromatic Features. In *Proceedings of the 2nd International Conference on Digital Tools* (Virtual Event, Tunisia) *(DTUC '20)*. Association for Computing Machinery, New York, NY, USA, Article 22, 6 pages. https://doi.org/10.1145/3423603.3424058

[11] Radhika Garg, Hua Cui, Spencer Seligson, Bo Zhang, Martin Porcheron, Leigh Clark, Benjamin R. Cowan, and Erin Beneteau. 2022. The Last Decade of HCI Research on Children and Voice-Based Conversational Agents. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 149, 19 pages. https://doi.org/10.1145/3491102.3502016

[12] Ben Hutchinson and Margaret Mitchell. 2019. 50 Years of Test (Un)Fairness: Lessons for Machine Learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) *(FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 49–58. https://doi.org/10.1145/3287560.3287600

[13] Kanae Kochigami, Jun Jiang, Yuriko Kakehashi, Chiwun Au, Yohei Kakiuchi, Kei Okada, and Masayuki Inaba. 2015. Walking together hand in hand: Design and evaluation of autonomous robot system that a robot recognizes moving direction with a child's assistance of pulling its hand. In *2015 IEEE/SICE International Symposium on System Integration (SII)*. 738–743. https://doi.org/10.1109/SII.2015.7405071

[14] N. Kumaresamoorthy and M.F.M. Firdhous. 2018. An APPROACH OF Filtering The Content Of Posts In Social Media. In *2018 3rd International Conference on Information Technology Research (ICITR)*. 1–6. https://doi.org/10.1109/ICITR.2018.8736152

[15] Christine P Lee, Bengisu Cagiltay, and Bilge Mutlu. 2022. The Unboxing Experience: Exploration and Design of Initial Interactions Between Children and Social Robots. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 151, 14 pages. https://doi.org/10.1145/3491102.3501955

[16] Kung Jin Lee, Wendy Roldan, Tian Qi Zhu, Harkiran Kaur Saluja, Sungmin Na, Britnie Chin, Yilin Zeng, Jin Ha Lee, and Jason Yip. 2021. The Show Must Go On: A Conceptual Model of Conducting Synchronous Participatory Design With Children Online. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 345, 16 pages. https://doi.org/10.1145/3411764.3445715

[17] Oussama Metatla, Sandra Bardot, Clare Cullen, Marcos Serrano, and Christophe Jouffrais. 2020. Robots for Inclusive Play: Co-Designing an Educational Game With Visually Impaired and Sighted Children. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376270

[18] Rebecca Nicholson, Tom Bartindale, Ahmed Kharrufa, David Kirk, and Caroline Walker-Gleaves. 2022. Participatory Design Goes to School: Co-Teaching as a Form of Co-Design for Educational Technology. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 150, 17 pages. https://doi.org/10.1145/3491102.3517667

[19] Netta Ofer, Idan David, Hadas Erel, and Oren Zuckerman. 2019. Coding for Outdoor Play: A Coding Platform for Children to Invent and Enhance Outdoor Play Experiences. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300394

[20] UK Parliament et al. 2022. Online Safety Bill (2022). (2022).

[21] Aditi Ramachandran, Chien-Ming Huang, and Brian Scassellati. 2017. Give Me a Break! Personalized Timing Strategies to Promote Learning in Robot-Child Tutoring. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna, Austria) *(HRI '17)*. Association for Computing Machinery, New York, NY, USA, 146–155. https://doi.org/10.1145/2909824.3020209

[22] Kaiwen Sun, Carlo Sugatan, Tanisha Afnan, Hayley Simon, Susan A. Gelman, Jenny Radesky, and Florian Schaub. 2021. "They See You're a Girl If You Pick a Pink Robot with a Skirt": A Qualitative Study of How Children Conceptualize Data Processing and Digital Privacy Risks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 687, 34 pages. https://doi.org/10.1145/3411764.3445333

[23] Kaiwen Sun, Yixin Zou, Jenny Radesky, Christopher Brooks, and Florian Schaub. 2021. Child Safety in the Smart Home: Parents' Perceptions, Needs, and Mitigation Strategies. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–41.

[24] Rashid Tahir, Faizan Ahmed, Hammas Saeed, Shiza Ali, Fareed Zaffar, and Christo Wilson. 2019. Bringing the Kid back into YouTube Kids: Detecting Inappropriate Content on Video Streaming Platforms. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. 464–469. https://doi.org/10.1145/3341161.3342913

[25] UNICEF. 2020. Policy guidance on AI for children. (2020).

[26] Ge Wang, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. 2021. Protection or Punishment? Relating the Design Space of Parental Control Apps and Perceptions about Them to Support Parenting for Online Safety. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 343 (oct 2021), 26 pages. https://doi.org/10.1145/3476084

[27] Julia Woodward, Feben Alemu, Natalia E. López Adames, Lisa Anthony, Jason C. Yip, and Jaime Ruiz. 2022. "It Would Be Cool to Get Stampeded by Dinosaurs": Analyzing Children's Conceptual Model of AR Headsets Through Co-Design. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 152, 13 pages. https://doi.org/10.1145/3491102.3501979

[28] Zheng Zhang, Ying Xu, Yanhao Wang, Bingsheng Yao, Daniel Ritchie, Tongshuang Wu, Mo Yu, Dakuo Wang, and Toby Jia-Jun Li. 2022. StoryBuddy: A Human-AI Collaborative Chatbot for Parent-Child Interactive Storytelling with Flexible Parental Involvement. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 218, 21 pages. https://doi.org/10.1145/3491102.3517479

[29] Jun Zhao, Ge Wang, Carys Dally, Petr Slovak, Julian Edbrooke-Childs, Max Van Kleek, and Nigel Shadbolt. 2019. 'I Make up a Silly Name': Understanding Children's Perception of Privacy Risks Online. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300336

[30] Oren Zuckerman, Guy Hoffman, Daphne Kopelman-Rubin, Anat Brunstein Klomek, Noa Shitrit, Yahav Amsalem, and Yaron Shlomi. 2016. KIP3: robotic companion as an external cue to students with ADHD. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. 621–626.