Beyond Users: Supporting Children in Interpreting, Resisting, and Collaborating with AI

Aayushi Dangol
Human Centered Design & Engineering
University of Washington
Seattle, Washington, USA
adango@uw.edu

Abstract

As Artificial Intelligence (AI) becomes increasingly embedded in children's everyday lives, the need to foster AI literacy from an early age has become more urgent. My dissertation will investigate how we can support children in engaging with AI not just as passive users, but as: 1) interpreters who make sense of AI's decision-making; 2) resistors who critically examine biased or flawed outputs; and 3) collaborators who co-create with AI in ways that reflect their identities, values, and lived experiences. To date, I have primarily explored the roles of interpreter and resistor through the co-design of interactive systems with children at KidsTeam UW, as well as through classroom deployments and field studies. This work has examined how children make sense of AI's decisions and how they critically respond when those decisions reflect bias or exclusion. Moving forward, I aim to investigate how AI can serve as a meaningful collaborator in supporting children's learning, development, and wellbeing, particularly in educational and therapeutic contexts.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**.

Keywords

AI Literacy, Generative AI

ACM Reference Format:
Aayushi Dangol. 2025. Beyond Users: Supporting Children in Interpreting, Resisting, and Collaborating with AI. In *Interaction Design and Children (IDC '25), June 23–26, 2025, Reykjavik, Iceland*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3713043.3731604

1 Introduction

Children today are growing up in a world where artificial intelligence (AI) is increasingly mediating how they interpret information, construct knowledge, and engage with the world around them. Recognizing the potential systemic and long-term impacts that AI systems may have on children's lives, scholars in AI education have suggested a series of essential competencies for navigating this technological landscape. Long and Magerko [30] define these competencies as the requisite skills for people to "*critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace*" [30, p. 2].

Over the past decade, while K-12 AI literacy initiatives have made significant strides in supporting these competencies, the rapid emergence of Large Language models (LLMs) and Large Reasoning models (LRMs) [16, 35] presents new challenges and opportunities for how children interact with AI. Unlike earlier AI systems, LLMs and LRMs are designed to generate human-like text, reason through complex problems, and adapt to diverse prompts with minimal explicit programming. Moreover, these models are now embedded in various applications [1, 4, 38, 51, 52], expanding children's interactions with AI beyond traditional embodied forms such as voice assistants and smart toys. However, these newer models are also prone to hallucinations and biased outputs [3, 24, 43], underscoring the importance of developing skills to interrogate model outputs, detect inconsistencies, and refine prompts for more accurate responses.

My research aims to explore how children understand and interact with emerging AI technologies, and the opportunities within these interactions to foster AI literacy. I draw inspiration from scholars in HCI, Learning Sciences, and AI Education to design child-centered systems that enable children to explore and engage in fluid experimentation with AI models. By designing learning experiences that help children connect with the ideas they encounter but also have opportunities to disconnect from them and reimagine new possibilities [40], my work seeks to empower children to adopt three key roles in their interactions with AI:

(1) **Interpreters**: Children develop foundational understanding of how AI processes data, makes decisions, and interprets the world.
(2) **Collaborators**: Children work alongside AI systems, leveraging its strengths while remaining aware of AI's limitations.
(3) **Resistors**: Children critically evaluate AI outputs, question its assumptions and, when necessary, reject AI outputs that may be flawed, biased, or ethically concerning.

2 Related Work

In this section, I first examine existing approaches to AI literacy in K-12 education, highlighting how different learning platforms have been designed to introduce children to AI concepts. Second, I explore prior work that integrates critical pedagogy into AI education, emphasizing how scholars have sought to help young learners recognize and challenge AI's socio-cultural implications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IDC '25, Reykjavik, Iceland
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1473-3/25/06
https://doi.org/10.1145/3713043.3731604

2.1 AI Literacy in K-12 Education

AI literacy encompasses an understanding of AI concepts, practices, and perspectives that enable learners to critically evaluate and utilize AI technologies [19, 30], while fostering considerations for AI ethics [30, 36, 42, 46]. Touretzky's five "big ideas" of AI – perception, representation & reasoning, learning, natural interaction, and societal impact – provide a strong foundation for fostering AI literacy [46]. In recent years, introducing young learners to AI concepts through accessible, hands-on experiences has become a central focus in AI literacy research [26, 34]. Over the past decade, researchers have developed a range of computational and unplugged learning platforms that help children understand AI's underlying mechanisms [25, 29, 32, 33, 44, 47].

Prior studies show that integrating open-ended platforms where children can experiment with AI models in a scaffolded and safe way helps them grasp AI concepts and engage in meaningful discussions about AI's capabilities and limitations [6, 20, 29, 39]. Platforms such as Teachable Machine [5], Poseblocks [25], and Cognimates [18] enable children to train AI models, observe predictions, and refine their AI models based on observed outcomes, making abstract AI concepts like classification and model prediction more tangible [2, 31, 37, 39, 48]. Research further indicates that aligning learning activities with children's personal experiences and interests increases engagement and helps build confidence among non-technical learners by creating a low barrier to entry [6]. At the same time, prior work also suggests that while uncovering some of these often "black-boxed" AI processes can improve understanding, revealing too much too quickly can overwhelm learners [22, 40]. Therefore, key aspects of AI concepts should be unveiled to children without overloading them. Building on these insights, I aim to design AI literacy experiences that inspire children to not only understand the underlying mechanisms of AI but also explore and experiment with it in ways that are personally meaningful to them.

2.2 Critical Inquiry of AI's Socio-Cultural Impact with Children

An early advocate of critical questioning in learning was Paulo Freire, whose work emphasizes the importance of teaching learners to question the systems that shape their lives and how they align with or challenge social justice and equity. Freire's call for learners to critically engage with "*the way they exist in the world with which and in which they find themselves*" [23, p.12] parallels the need in AI education to examine how the values embedded in AI systems interact with societal structures, potentially reinforcing inequities or biases.

Recently, various educational resources aimed at different age groups have been introduced to raise awareness of AI bias and fairness at the K-12 level. For example, Melsión et al. [32] developed an educational platform to highlight gender bias in supervised machine learning, encouraging pre-adolescents to critically evaluate how bias can influence algorithmic decisions. Similarly, Coenraad et al. [9] explored AI bias within the broader context of technology's impact on youth, demonstrating how young people could recognize and provide examples of bias in their own lives. Other scholars have approached bias in AI technologies within the context of teaching children about algorithmic fairness [21, 27, 41], and have demonstrated that children are capable of recognizing bias within their lives and technologies [9, 21, 27, 41]. By building on this intuitive understanding of bias and fairness in children's lives, prior work has highlighted how children's existing perceptions of bias can inform a deeper understanding of AI systems and its ethical implications. My work builds on this approach by supporting children in understanding the bidirectional relationship between AI and society [49], where AI is molded by society's values, and, in turn, society is influenced by AI.

3 Contributions To Date

The following sections outline my contributions to designing AI systems that support children's understanding of AI, while empowering them to critique and reshape its role in their lives. This work builds on my earlier explorations of data literacy [7, 10], which examined how children interpret data and interrogate the assumptions embedded within it.

3.1 Supporting Children as Interpreters of AI Decision-Making

A central thread in my research explores how children interpret AI's decision-making, what kinds of reasoning they attribute to AI, and how their understandings evolve through actively probing and manipulating AI behaviors. In Study 1, I conducted a two-part investigation to examine how children conceptualize AI reasoning, using Abstractions & Reasoning Corpus [8], a benchmark for evaluating abstract reasoning in AI systems, as a developmentally appropriate scaffold. Findings from a co-design session with 8 children followed by a field study with 106 children (grades 3–8) revealed three distinct mental models: Inherent (*AI reasons because it is inherently intelligent*), Deductive (*AI follows rules it was programmed with*), and Inductive (*AI generalizes from data to recognize patterns*) [14]. Additionally, we found a developmental shift where younger children often viewed AI as intrinsically intelligent, while older children were more likely to describe AI as a pattern recognizer. These findings highlight both children's intuitive insights and the misconceptions they carry, underscoring the need for educational tools that scaffold deeper engagement with how AI reasons.

Building on this foundational work, in Study 2, I developed AI-Smartlock, a system that simulates a smart door lock to help children understand classification-based AI reasoning.[1] The system introduces three classification parameters: classification rules (categorizing data), confidence scores (estimating classification certainty), and decision thresholds (setting minimum confidence levels for acceptance). Findings from a co-design session with seven children (ages 7–12) revealed that open-ended exploration with AI-Smartlock helped children compare their own decision-making processes with those of the AI. Through iterative experimentation with classification parameters, children formed hypotheses about AI's decision-making, tested their predictions, and refined their understanding of AI classification.

While AI-Smartlock focused on how children interpret AI's classification logic, my second study extends this inquiry to the more ambiguous and often unpredictable domain of generative AI. In

[1]Accepted to *The International Conference of the Learning Sciences (ISLS 2025).*

Study 2, I developed AI Puzzlers, a web-based game that helps children explore generative AI's reasoning by solving visual puzzles [15]. Findings from co-design sessions with 21 children (ages 6 - 11) suggest that when generative AI made mistakes on seemingly simple visual puzzles, children's surprise led to deeper inquiry into how and why AI arrives at certain conclusions. Even younger children, who were not yet fluent readers, quickly detected inconsistencies in AI-generated solutions, prompting discussions about why AI's reasoning differs from human logic. Children's hands-on engagement with the system made AI reasoning more accessible by giving them a concrete way to observe and question AI's decisions. It also helped them recognize both generative AI's capabilities and its limitations. By shifting the focus from understanding AI as a "black box" to actively probing and manipulating AI behaviors, my research contributes to ongoing discussions in HCI, AI literacy, and child-centered AI design. In my future work, I hope to explore how these insights can inform AI education curricula and the development of explainable AI interfaces that better support children's learning about AI decision-making.

3.2 Supporting Children as Resistors of Biased and Flawed AI Outputs

My work aims to equip children with the tools to critically engage with AI technologies, empowering them to play a meaningful role in shaping ethical and responsible AI systems. I approach this goal through participatory design and critical pedagogy. In Study 3, I investigated how societal attitudes toward adolescents are reflected in generative Large Language Models (LLMs), comparing these computational representations to the lived perspectives of adolescents themselves [50]. Recognizing that attitudes toward youth vary across cultures, I implemented a bilingual, bicultural study with 13 English-speaking adolescents in the U.S. and 18 Nepali-speaking adolescents in Nepal. In these sessions, young people (ages 13 - 17) reflected on how teenagers are portrayed in digital media and discussed how they want to be represented by AI systems. Findings show that LLMs often sensationalize adolescence – emphasizing violence, drug use, or mental health crises – while overlooking the everyday richness and diversity of young people's lives. Importantly, the study revealed adolescents' desire to participate in shaping how they are represented in AI systems and their capacity to meaningfully engage in conversations about data, design, and fairness.

Extending this exploration of AI's role in shaping cultural narratives and representations, in Study 4, I used participatory design to explore how a group of 13 children (ages 8 - 13) understand generative AI's role in mediating culture [13]. The findings revealed that children are developing a nuanced awareness of how AI systems are influenced by, and in turn influence, cultural data. They identified training data as a key limitation in AI-mediated culture, recognizing that the selection of training data, algorithmic rules, and fine-tuning processes all play a critical role in how AI constructs and conveys cultural concepts.

Lastly, in Study 5, I expand on the theme of critical engagement by connecting children's understanding of AI bias to broader societal structures of inequality.[2] Building on insights from my previous studies, I designed and implemented an interactive system called CLIP4KIDS with 28 fifth-grade students (ages 9 - 10) at a local school in the Pacific Northwest, United States. Through hands-on engagement with the system and guided dialogue facilitated by their technology teacher, the students actively reflected on and critiqued flawed and biased AI outputs. By engaging in these discussions, children demonstrated an awareness that AI is not separate from the societal contexts in which it is created, but is shaped by those very contexts. Across these three studies, AI emerges as a site of negotiation and resistance—a space where children don't just passively navigate bias, but actively challenge, re-imagine, and construct alternative technological futures. By centering their voices in AI design, my work envisions AI not as a force that defines children, but as a system that can be co-created with them to better reflect their diverse experiences and identities.

[2]Under review.

4 Current Work and Next Steps

In my future work, I aim to explore how AI can act as a collaborator in supporting children's development and wellbeing, particularly within educational and therapeutic settings [11, 17, 28, 45]. I have begun this exploration through Study 6, which investigates how AI might assist parents in supporting their children's speech therapy at home [12]. This study involved semi-structured interviews with 20 parents of children with speech difficulties, focusing on the practical and emotional challenges of home practice and how AI could help address them. Building on this, I am now working with eight families to better understand how children's exploratory, multisensory interaction styles during speech practice might inform the design of supportive AI technologies. Through observations, cultural probes, and follow-up interviews with caregivers, I aim to identify moments in daily routines where speech practice can naturally align with children's interests and activities. Additionally, I seek to understand how caregivers scaffold learning through storytelling, play, and other interactive strategies, and explore how AI tools might complement these natural, relational practices.

Acknowledgments

This material is based upon work supported under the AI Research Institutes program by the National Science Foundation and the Institute of Education Sciences, U.S. Department of Education, through Award #DRL-2229873 - AI Institute for Transforming Education for Children with Speech and Language Processing Challenges (or National AI Institute for Exceptional Education). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, the Institute of Education Sciences, or the U.S. Department of Education. This work was also partially funded by the Jacob's Foundation CERES Network.

References

[1] Eman A Alasadi and Carlos R Baiz. 2023. Generative AI in education and research: Opportunities, concerns, and solutions. *Journal of Chemical Education* 100, 8 (2023), 2965–2971.
[2] Safinah Ali, Daniella DiPaola, Irene Lee, Victor Sindato, Grace Kim, Ryan Blumofe, and Cynthia Breazeal. 2021. Children as Creators, Thinkers and Citizens in an AI-driven Future. 2 (2021), 100040. doi:10.1016/j.caeai.2021.100040
[3] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models

be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.

[4] Nadia Bidarian. 2023. Meet Khan Academy's Chatbot Tutor | CNN business. https://www.cnn.com/2023/08/21/tech/khan-academy-ai-tutor/index.html

[5] Michelle Carney, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. Teachable machine: Approachable Web-based tool for exploring machine learning classification. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*. 1–8.

[6] Francisco Enrique Vicente Castro, Kayla DesPortes, William Payne, Yoav Bergner, and Kathleen McDermott. 2022. AI+ Dance: Co-Designing Culturally Sustaining Curricular Resources for AI and Ethics Education Through Artistic Computing. In *Proceedings of the 2022 ACM Conference on International Computing Education Research-Volume 2*. 26–27.

[7] Ruijia Cheng, Aayushi Dangol, Frances Marie Tabio Ello, Lingyu Wang, and Sayamindu Dasgupta. 2023. Concepts, practices, and perspectives for developing computational data literacy: Insights from workshops with a new data programming system. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference*. 100–111.

[8] François Chollet. 2019. On the measure of intelligence. *arXiv preprint arXiv:1911.01547* (2019).

[9] Merijke Coenraad. 2022. "That's what techquity is": youth perceptions of technological and algorithmic bias. *Information and Learning Sciences* 123, 7/8 (Jan. 2022), 500–525. doi:10.1108/ILS-03-2022-0024 Publisher: Emerald Publishing Limited.

[10] Aayushi Dangol and Sayamindu Dasgupta. 2023. Constructionist approaches to critical data literacy: A review. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference*. 112–123.

[11] Aayushi Dangol, Yun Huang, Srirangaraj Setlur, Adele Smolansky, Hariharan Subramonyam, Hyewon Suh, Jinjun Xiong, and Julie A Kientz. 2024. AI-Driven Support for People with Speech & Language Difficulties. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–4.

[12] Aayushi Dangol, Aaleyah Lewis, Hyewon Suh, Xuesi Hong, Hedda Meadan, James Fogarty, and Julie A Kientz. 2025. "I Want to Think Like an SLP": A Design Exploration of AI-Supported Home Practice in Speech Therapy. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21. doi:10.1145/3706598.3713986

[13] Aayushi Dangol, Michele Newman, Robert Wolfe, Jin Ha Lee, Julie A Kientz, Jason Yip, and Caroline Pitt. 2024. Mediating Culture: Cultivating Socio-cultural Understanding of AI in Children through Participatory Design. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 1805–1822.

[14] Aayushi Dangol, Robert Wolfe, Runhua Zhao, JaeWon Kim, Trushaa Ramanan, Katie Davis, and Julie A. Kientz. 2025. Children's Mental Models of AI Reasoning: Implications for AI Literacy Education. In *Proceedings of the Interaction Design and Children Conference (IDC '25)*. Association for Computing Machinery, Reykjavik, Iceland, 18. doi:10.1145/3713043.3728856

[15] Aayushi Dangol, Runhua Zhao, Robert Wolfe, Trushaa Ramanan, Julie A. Kientz, and Jason Yip. 2025. "AI just keeps guessing": Using ARC Puzzles to Help Children Identify Reasoning Errors in Generative AI. In *Proceedings of the Interaction Design and Children Conference (IDC '25)*. Association for Computing Machinery, Reykjavik, Iceland. doi:10.1145/3713043.3728836

[16] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948 [cs.CL] https://arxiv.org/abs/2501.12948

[17] Riddhi Divanji, Aayushi Dangol, Ella J Lombard, Katharine Chen, and Jennifer D Rubin. 2024. Togethertales RPG: Prosocial skill development through digitally mediated collaborative role-playing. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference*. 1012–1015.

[18] Stefania Druga. 2018. *Growing up with AI: Cognimates: from coding to teaching machines*. Ph. D. Dissertation. Massachusetts Institute of Technology.

[19] Stefania Druga, Fee Lia Christoph, and Amy J Ko. 2022. Family as a Third Space for AI Literacies: How do children and parents learn about AI together?. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.

[20] Stefania Druga and Amy J Ko. 2021. How do children's perceptions of machine intelligence change when training and coding smart programs?. In *Interaction Design and Children*. ACM, Athens Greece, 49–61. doi:10.1145/3459990.3460712

[21] Stefania Druga, Jason Yip, Michael Preston, and Devin Dillon. 2021. The 4As: Ask, adapt, author, analyze-AI literacy framework for families. *Algorithmic Rights and Protections for Children* (2021).

[22] Utkarsh Dwivedi, Jaina Gandhi, Raj Parikh, Merijke Coenraad, Elizabeth Bonsignore, and Hernisa Kacorri. 2021. Exploring Machine Teaching with Children. In *2021 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, Vol. 2021. IEEE, United States, 1–11.

[23] Paulo Freire. 1978. Pedagogy of the oppressed. In *Toward a Sociology of Education*. Routledge, 374–386.

[24] Wei Guo and Aylin Caliskan. 2021. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 122–133.

[25] Brian Jordan, Nisha Devasia, Jenna Hong, Randi Williams, and Cynthia Breazeal. 2021. PoseBlocks: A toolkit for creating (and dancing) with AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 15551–15559.

[26] Ken Kahn and Niall Winters. 2021. Constructionism and AI: A history and possible futures. *British Journal of Educational Technology* 52, 3 (2021), 1130–1142.

[27] Victor R. Lee, Victoria Delaney, and Parth Sarin. 2022. Eliciting High School Students' Conceptions and Intuitions about Algorithmic Bias. In *Proceedings of the 2022 ACM Conference on International Computing Education Research - Volume 2* (Lugano and Virtual Event, Switzerland) *(ICER '22)*. ACM, New York, NY, USA, 35–36. doi:10.1145/3501709.3544279

[28] Aaleyah Lewis, Aayushi Dangol, Hyewon Suh, Abbie Olszewski, James Fogarty, and Julie A. Kientz. 2025. Exploring AI-Based Support in Speech-Language Pathology for Culturally and Linguistically Diverse Children. In *CHI Conference on Human Factors in Computing Systems (CHI '25)* (Yokohama, Japan). Association for Computing Machinery, ACM, New York, NY, USA. doi:10.1145/3706598.3714131

[29] Duri Long, Takeria Blunt, and Brian Magerko. 2021. Co-Designing AI Literacy Exhibits for Informal Learning Spaces. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–35. doi:10.1145/3476034

[30] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. doi:10.1145/3313831.3376727

[31] Radu Mariescu-Istodor and Ilkka Jormanainen. 2019. Machine learning for high school students. In *Proceedings of the 19th Koli calling international conference on computing education research*. 1–9.

[32] Gaspar Isaac Melsión, Ilaria Torre, Eva Vidal, and Iolanda Leite. 2021. Using Explainability to Help Children Understand Gender Bias in AI. In *Interaction Design and Children (IDC '21)*. ACM, New York, NY, USA, 87–99. doi:10.1145/3459990.3460719

[33] Terran Mott, Alexandra Bejarano, and Tom Williams. 2022. Robot co-design can help us engage child stakeholders in ethical reflection. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 14–23.

[34] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence* 2 (2021), 100041.

[35] OpenAI. 2024. Learning to reason with LLMs. *OpenAI Blog* , (Sep 2024), .

[36] Blakeley H Payne. 2019. An ethics of artificial intelligence curriculum for middle school students. *MIT Media Lab Personal Robots Group. Retrieved Oct* 10 (2019), 2019.

[37] Blakeley H Payne. 2020. *Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students*. Ph. D. Dissertation. Massachusetts Institute of Technology.

[38] Junaid Qadir. 2023. Engineering education in the era of ChatGPT: Promise and pitfalls of generative AI for education. In *2023 IEEE Global Engineering Education Conference (EDUCON)*. IEEE, 1–9.

[39] Yim Register and Amy J Ko. 2020. Learning machine learning with personal data helps stakeholders ground advocacy arguments in model mechanics. In *Proceedings of the 2020 ACM Conference on International Computing Education Research*. 67–78.

[40] Mitchel Resnick and Eric Rosenbaum. 2013. Designing for tinkerability. *Design, make, play: Growing the next generation of STEM innovators* (2013), 163–181.

[41] Jean Salac, Alannah Oleson, and Lena Armstrong. 2023. Funds of Knowledge used by Adolescents of Color in Scaffolded Sensemaking around Algorithmic Fairness. (2023).

[42] Jeffrey Saltz, Michael Skirpan, Casey Fiesler, Micha Gorelick, Tom Yeh, Robert Heckman, Neil Dewar, and Nathan Beard. 2019. Integrating ethics within machine learning courses. *ACM Transactions on Computing Education (TOCE)* 19, 4 (2019), 1–26.

[43] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2019. The Woman Worked as a Babysitter: On Biases in Language Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3407–3412.

[44] Jaemarie Solyst, Amy Ogan, and Jessica Hammer. 2023. Intergenerational Games to Learn About AI and Ethics. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 2* (Toronto ON, Canada) *(SIGCSE 2023)*. Association for Computing Machinery, New York, NY, USA, 1273. doi:10.1145/3545947.3573256

[45] Hyewon Suh, Aayushi Dangol, Hedda Meadan, Carol A Miller, and Julie A Kientz. 2024. Opportunities and Challenges for AI-Based Support for Speech-Language Pathologists. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. Association for Computing Machinery, Newcastle upon Tyne, United Kingdom, 1–14.

[46] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What should every child know about AI?. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 9795–9799.

[47] Jessica Van Brummelen, Viktoriya Tabunshchyk, and Tommy Heng. 2021. "Alexa, Can I Program You?": Student Perceptions of Conversational Artificial Intelligence Before and After Programming Alexa. In *Proceedings of the 20th Annual ACM Interaction Design and Children Conference* (New York, NY, USA, 2021-06-24) *(IDC '21)*. Association for Computing Machinery, 305–313. doi:10.1145/3459990.3460730

[48] Iro Voulgari, Marvin Zammit, Elias Stouraitis, Antonios Liapis, and Georgios Yannakakis. 2021. Learn to machine learn: designing a game based approach for teaching machine learning to primary and secondary education students. In *Interaction design and children*. 593–598.

[49] Robert Wolfe, Aayushi Dangol, Alexis Hiniker, and Bill Howe. 2024. Dataset Scale and Societal Consistency Mediate Facial Impression Bias in Vision-Language AI. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, Vol. 7. 1635–1647.

[50] Robert Wolfe, Aayushi Dangol, Bill Howe, and Alexis Hiniker. 2024. Representation Bias of Adolescents in AI: A Bilingual, Bicultural Study. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, Vol. 7. 1621–1634.

[51] Chao Zhang, Cheng Yao, Jianhui Liu, Zili Zhou, Weilin Zhang, Lijuan Liu, Fangtian Ying, Yijun Zhao, and Guanyun Wang. 2021. StoryDrawer: A Co-Creative Agent Supporting Children's Storytelling through Collaborative Drawing. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–6.

[52] Yijun Zhao, Yiming Cheng, Shiying Ding, Yan Fang, Wei Cao, Ke Liu, and Jiacheng Cao. 2024. Magic Camera: An AI Drawing Game Supporting Instantaneous Story Creation for Children. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference*. 738–743.