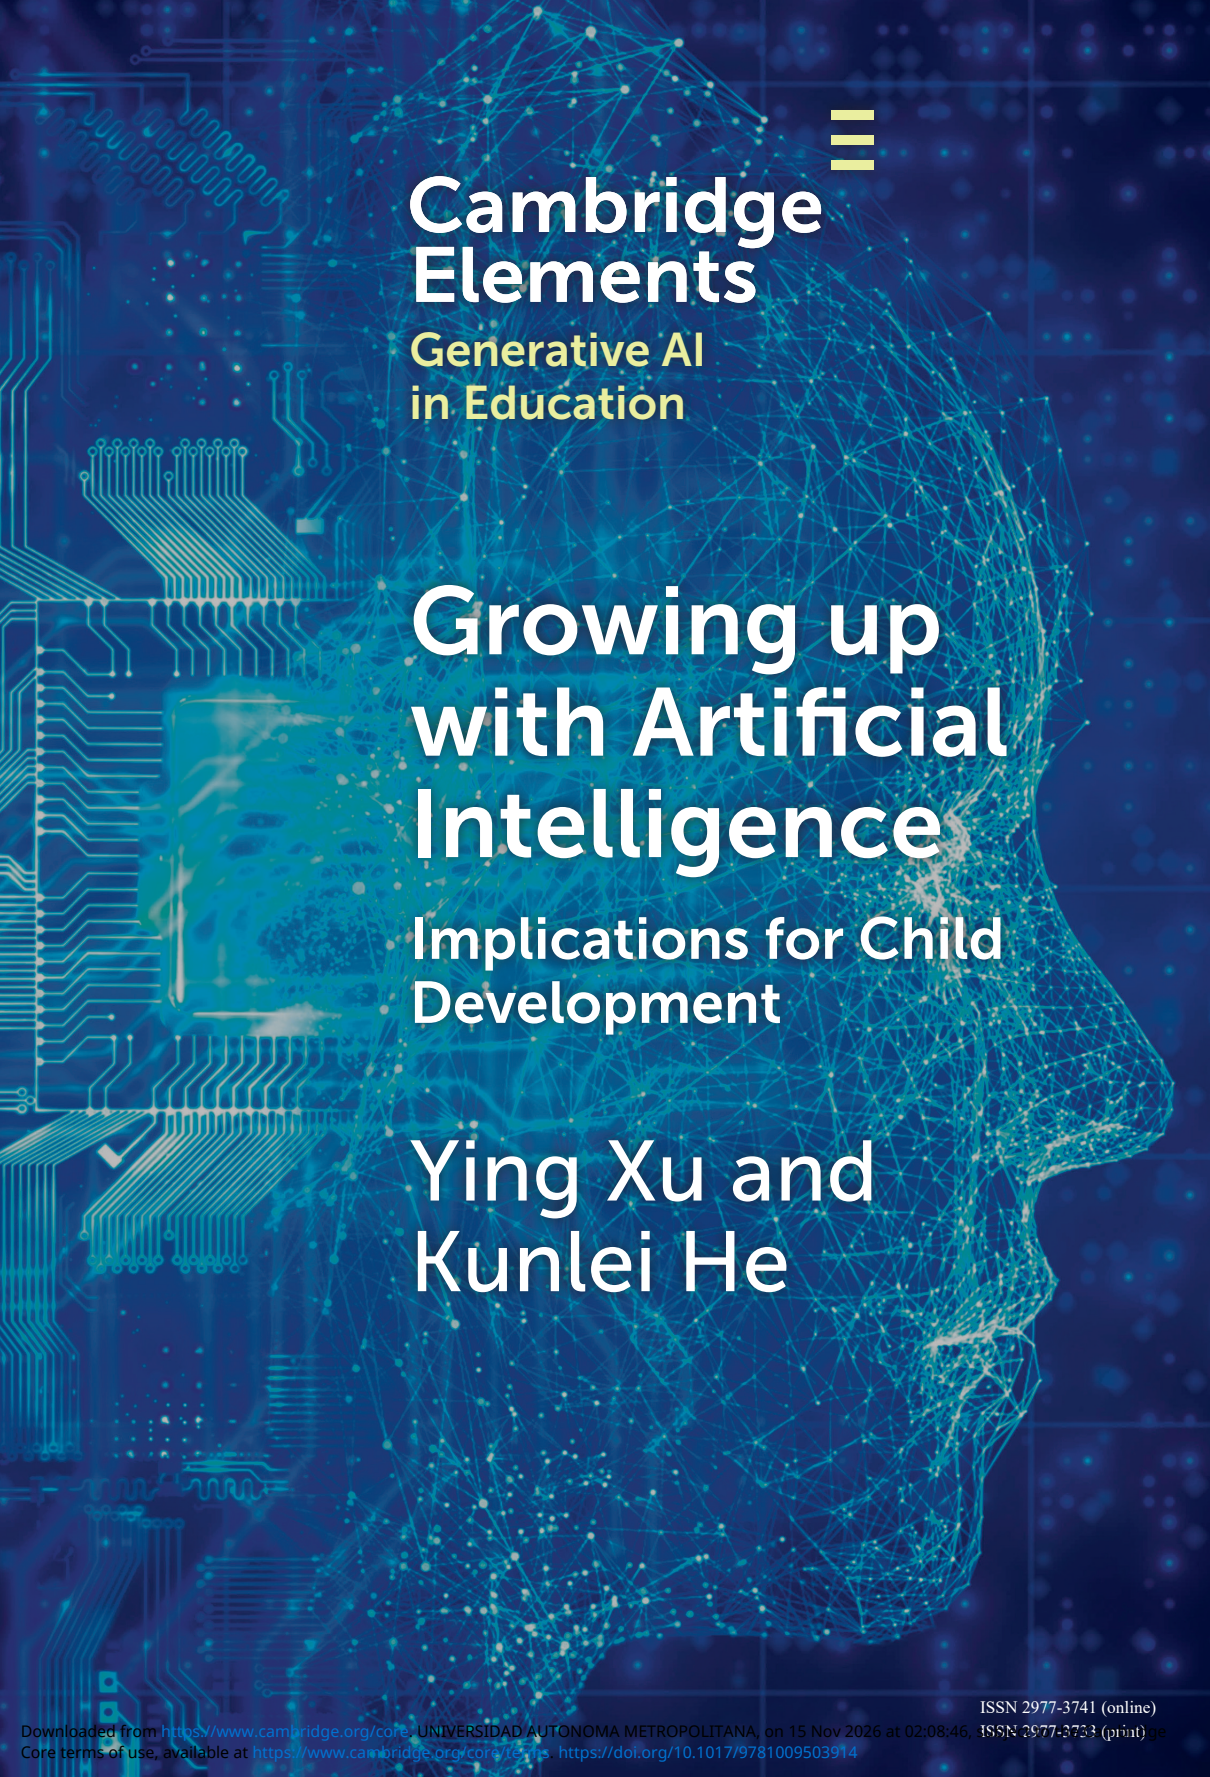


Cambridge Elements

Generative AI in Education

Growing up with Artificial Intelligence

Implications for Child Development

Ying Xu and Kunlei He

ISSN 2977-3741 (online)
ISSN 2977-3733 (print)

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Cambridge Elements

Elements in Generative AI in Education

edited by

Mark Warschauer

University of California, Irvine

Tamara Tate

University of California, Irvine

GROWING UP WITH ARTIFICIAL INTELLIGENCE

Implications for Child Development

Ying Xu

Harvard University

Kunlei He

University of California, Irvine

CAMBRIDGE UNIVERSITY PRESS

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

CAMBRIDGE UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781009503907

DOI: 10.1017/9781009503914

© Ying Xu and Kunlei He 2026

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, with the exception of the Creative Commons version the link for which is provided below, no reproduction of any part may take place without the written permission of Cambridge University Press & Assessment.

An online version of this work is published at doi.org/10.1017/9781009503914 under a Creative Commons Open Access license CC-BY-NC-ND 4.0 which permits re-use, distribution and reproduction in any medium for non-commercial purposes providing appropriate credit to the original work is given. You may not distribute derivative works without permission. To view a copy of this license, visit https://creativecommons.org/licenses/by-nc-nd/4.0

When citing this work, please include a reference to the DOI 10.1017/9781009503914

First published 2026

A catalogue record for this publication is available from the British Library

A Cataloging-in-Publication data record for this Element is available from the Library of Congress

ISBN 978-1-009-50390-7 Hardback
ISBN 978-1-009-50386-0 Paperback
ISSN 2977-3741 (online)
ISSN 2977-3733 (print)

Cambridge University Press & Assessment has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

For EU product safety concerns, contact us at Calle de José Abascal, 56, 1°, 28003 Madrid, Spain, or email eugpsr@cambridge.org

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Growing up with Artificial Intelligence

Implications for Child Development

Elements in Generative AI in Education

DOI: 10.1017/9781009503914
First published online: May 2026

Ying Xu
Harvard University

Kunlei He
University of California, Irvine

Author for correspondence: Ying Xu, yingxu@g.harvard.edu

Abstract: Generative artificial intelligence (AI) is becoming an integral part of children's lives, ranging from voice assistants and social robots to AI-generated storybooks. As children increasingly interact with these technologies, it is essential to consider their implications for developmental outcomes. This Element examines these implications across three interconnected domains: interaction, perception, and learning. A recurring theme across these domains is that children's engagement with AI both mirrors and diverges from their engagement with humans, positioning AI as a distinct yet potentially complementary source of experience, enrichment, and knowledge. Ultimately, the Element advances a framework for understanding the complex interplay among technology, children, and the social contexts that shape their development. This title is also available as Open Access on Cambridge Core.

Keywords: artificial intelligence, children, cognitive development, communication, social development

© Ying Xu and Kunlei He 2026

ISBNs: 9781009503907 (HB), 9781009503860 (PB), 9781009503914 (OC)
ISSNs: 2977-3741 (online), 2977-3733 (print)

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Contents

1	Introduction	1
2	A Note on Research Methods	4
3	Children's Interactions with AI Agents	8
4	Children's Perceptions of AI	22
5	Children's Trust in AI-Generated Information	28
6	AI and Children's Learning	33
7	Children's Learning about AI	45
8	Conclusion	49
	References	53

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

1 Introduction

Artificial intelligence (AI) has become deeply integrated into modern society, and many children today are regularly exposed to AI-powered devices, content, and experiences. According to a recent national survey, an overwhelming majority (87 percent) of American children aged three to twelve now have access to AI-powered devices, from voice assistants to smartphones, with half using them daily (Bickham et al., 2024). Children increasingly rely on digital assistants like Siri, Alexa, or ChatGPT to ask questions and seek information, engage in interactive storytelling with social robots, or play with smart toys that can recognize and respond to facial expressions. Additionally, AI algorithms are now integral to adaptive video games and intelligent tutoring systems, which personalize content to suit individual learning needs. More recently, AI has entered the children's publishing industry, producing storybooks with polished narratives and AI-generated illustrations. The growing presence of AI has sparked significant attention – both in academia and among the general public – about the potential impact of AI on children's development.

The term "AI" encompasses a wide range of technologies, including computer vision, classification, robotics, natural language processing, and speech recognition. This breadth is reflected in the following definition by the Organisation for Economic Co-operation and Development (OECD): "A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. It uses machine and/or human-based inputs to perceive real and/or virtual environments, abstracts such perceptions into models, and employs model inference to formulate options for information or action" (OECD, 2019). While AI is a broad and complex concept, it is often treated in public discourse and everyday understanding as a singular unified technology. For the purpose of this Element, to avoid both oversimplification and getting too caught up in technical details, we adopt a more pragmatic approach to define what we mean by AI rather than relying solely on a technological perspective. We use the term "AI" to refer to the group of technologies or devices that children perceive and engage with as interactive partners – often known as "conversational agents" – which include voice assistants (e.g., Siri, Alexa), chatbots, AI-enabled toys, and social robots. While the underlying technologies that power these agents may vary – often combining multiple methods such as natural language processing, speech recognition, and machine learning – their defining characteristic is the capacity to support meaningful, human-like communication between users and machines. This kind of conversational AI is particularly relevant to child development because conversation plays a crucial role in shaping their growth (Xu, 2023). Decades of research has demonstrated that engaging in back-and-forth dialogues

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

with parents, siblings, teachers, and peers helps children develop language skills and deepen their understanding of others and the world around them. Both the quantity and quality of language exposure significantly influence various developmental outcomes (Golinkoff et al., 2019). Furthermore, conversations are vital for children to learn concepts and acquire knowledge they might not otherwise encounter firsthand (Gelman, 2003). This raises important questions about how interactions with AI fit within, and potentially reshape, the broader landscape of children's conversational experiences.

It is fair to say that many of our discussions and concerns about AI in relation to children are not entirely new. Since the introduction of digital media technologies for children – beginning with earlier media like television, followed by mobile apps and, more recently, social media – there has been ongoing debate about how these tools influence the way children engage in learning experiences and interact with others, and how this may affect their development. Some are concerned about "displacement" – the idea that excessive use of these technologies by children could reduce the amount and quality of other important activities, thus potentially hindering their social and cognitive development. On the other hand, some believe these technologies offer new ways to enhance and enrich the way children learn. For example, some studies have shown that when parents and children interact with digital books together, children may focus more on the digital features, which can lead to less meaningful conversation (e.g., Munzer et al., 2019). At the same time, digital books with narration features help preliterate children enjoy stories on their own when caregivers are not available (Egert et al., 2022).

We face similar questions when it comes to conversational AI, which could have both positive and negative implications for children's development. On the one hand, AI can greatly expand children's access to information, making it easier for them to obtain knowledge by simply asking, much like they would from other people. It also makes it more feasible to cultivate personalized learning experiences for children. On the other hand, AI's ability to provide quick answers might, to some extent, reduce the need for children to engage in problem-solving and critical reflection, potentially limiting their opportunities to develop both critical thinking and foundational academic skills. In addition, while AI portrayed as a companion might offer children a sense of companionship when desired, there are concerns about how this might affect their relationships with the people around them. These questions about whether conversational AI benefits or hinders children's learning and development fuel the polarized views that drive much of the current discussion. These views often remain speculative and opinion based due to the lack of direct evidence or the inconclusive nature of the available evidence thus far.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

This Element aims to provide an overview of what is currently known about children's development in the era of AI. At the same time, it acknowledges the many different questions that remain unanswered and seeks to foster readers' curiosity about these unknowns. To bring structure to this rapidly emerging landscape, the Element focuses on three central areas: interactions, perceptions, and learning. The first area focuses on how interactions serve as a foundation for how children engage with AI. These interactions shape the relationships children form with AI and set the stage for their understanding and learning. The second area, perceptions and beliefs, examines how children make sense of AI based on their experiences and existing understanding of the world. Their perceptions are influenced both by their direct experiences and knowledge of AI and their developmental capacities. The final area, learning and development, explores how AI affects children's growth in areas like language development, subject-domain learning, problem-solving, and creativity. It also considers the potential benefits and risks of AI in these areas, emphasizing the need for thoughtful design and use of these technologies. Figure 1 illustrates the overall organization of the Element and the relationships among the three areas. We will examine these relationships in greater detail in the Conclusion.

When we use the term "children" in this Element, we are primarily referring to those in early childhood – specifically, children in preschool through the early elementary years. This focus is deliberate for two interrelated reasons. First, our inquiry into AI's implications for child development mirrors long-standing questions about how children acquire knowledge, form beliefs, and interact with social agents – whether human or artificial. This inquiry is closely linked to the following question: To what extent do children distinguish between humans and machines? This question is deeply tied to their developing abilities in categorization and understanding social groupings (Rakison & Dubois, 2001). Early theorists like Jean Piaget explained that young children often use "magical thinking" or anthropomorphism – that is, they may believe nonhuman things have human-like qualities. Yet contemporary psychologists argue that children's classifications are not based solely

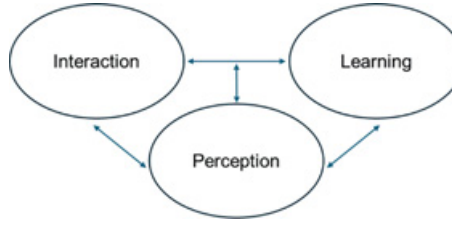


Figure 1 Three focal components concerning children's interactions with AI.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

on salient appearances but also on underlying essences – stable, internal properties that define category membership – showing that children have a more sophisticated understanding than previously thought (Gelman, 2003). Children from preschool through elementary school begin to form foundational beliefs about the nature of things, which influence how they categorize and make sense of the world around them. Studying children's interactions with AI during this period provides an opportunity to explore how these early concepts of category and agency develop when children encounter machines that can behave in human-like ways. Second, and relatedly, for many young children, interactions with AI – such as smart toys, voice assistants, or learning platforms – represent their first encounters with intelligent systems. These early experiences often occur before the influence of formal schooling or strong exposure to cultural narratives about what AI is, offering a unique window into their intuitive and spontaneous thinking.

2 A Note on Research Methods

Exploring the intersection of child development and AI is inherently interdisciplinary, drawing from psychology, computer science, communication, and education – each with its own rich history of research. The current state of research reflects this interdisciplinary nature, grounded in different theories, methodologies, and approaches. Inevitably, these methodologies come with their own strengths and weaknesses, as well as different ways in which the results should be interpreted. Thus, it is necessary to briefly clarify the common methodologies used by researchers to examine AI and child development before diving into the specific study results.

2.1 Types of AI Used in Research Studies

Researchers have used AI in very different ways across studies – some rely on commercial products like Alexa, others use custom-built research prototypes, and some use simulated systems that are actually controlled by humans. These choices affect how we interpret the findings. For example, they influence how generalizable the results are, whether the outcomes could be replicated in real-world settings, and whether the findings reflect typical experiences or ideal, best-case scenarios.

The first type involves using commercially available products. Some of these are specifically designed for children, such as the Echo Dot Kids Edition and Khanmigo, while others are more general purpose and not tailored to child users such as ChatGPT. Given that these products are readily available and commonly used by children, studies using them typically yield results with strong

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

ecological validity, as they reflect real-world conditions and naturalistic usage patterns. For instance, according to Khanmigo's report (Khan Academy, 2024), more than 200,000 individuals used their product during the 2023–2024 school year. Students often use Khanmigo as a tutor, seeking step-by-step support when they get stuck. However, a limitation of using off-the-shelf systems for research is that, since these products are already developed and encapsulated in commercial software, researchers have limited flexibility to modify them. This makes it challenging to isolate and examine the specific mechanisms that might influence children's interactions with these systems or the learning outcomes they produce. In addition, some of the products studied might not be specifically designed for children (e.g., ChatGPT); as a result, the affordances or limitations observed in these studies might not represent the most ideal outcomes AI could bring to children.

The second type of AI systems used in studies consists of proof-of-concept prototypes or research-developed tools that are intentionally designed to promote specific learning or developmental outcomes. These systems are often grounded in particular theories of child development or pedagogy and are built to embody strategies believed to support those outcomes – such as fostering self-regulation, encouraging dialogic learning, or scaffolding problem-solving. This intentional alignment between theory and design allows researchers to isolate and test the impact of specific features. By comparing children's outcomes with and without these targeted features or different design principles, studies can generate precise insights into which design elements are effective and under what conditions. As a result, this type of research is especially valuable for advancing both theory and practice – it helps identify not just *whether* an AI tool works but *why* it works.

Lastly, the third type of AI is "simulated AI." This method, often called the "Wizard of Oz" approach, involves researchers secretly manually controlling devices while informing participants that they are interacting with an AI tool. The key advantage of this approach is that it allows the research team to precisely control the AI tool's behaviors, avoiding confounding factors caused by unpredictable AI errors or limitations. For example, researchers typically follow a strict, predefined script when operating the simulated AI tool, which may even include staged errors to observe specific reactions or scenarios. Moreover, although this method does not employ autonomous AI systems, the rapid advancement of AI technologies toward human-like behaviors makes this approach increasingly relevant. Indeed, using human operators to simulate AI can provide a valuable, forward-looking perspective on how AI may function in the near future.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

2.2 Study Design

In addition to understanding the type of AI implementation, it is also important to consider how studies on AI are designed, as this shapes the kinds of questions researchers can ask and the conclusions they can draw. Broadly, studies examining AI and children usually use three main types of design.

The first type comprises descriptive studies, which aim to document and understand patterns of behavior. Researchers use tools like surveys, interviews, field notes, and fine-grained log data to capture children's access to or engagement with AI, which helps to identify trends and describe phenomena. Several notable examples include recent pulse surveys on children's use of AI tools (Bickham et al., 2024; Madden et al., 2024); they offer readers a bird's-eye view of the rapid proliferation of AI adoption in early childhood. Other descriptive studies may be smaller in scale but use more in-depth data, such as studies that use log data to record the interactions of a few dozen children with AI over time and understand the types of questions they initiate (Oh et al., 2025).

The second type of design consists of correlational studies, which examine relationships between variables to explore how certain factors may predict or relate to children's interactions with or learning from AI. For example, researchers might investigate whether a child's age, language background, or prior technology exposure predicts how often they engage with an AI tool or how much they learn from it. While correlational studies do not prove causation, they are valuable for identifying meaningful patterns – such as whether children from higher-income households are more likely to access AI-driven educational apps, or whether children who have more frequent interactions with AI systems show stronger gains in certain skills (Klarin et al., 2024). These findings can point researchers toward important questions to test further through experiments.

The third type of design comprises experimental studies, which are designed to test causal relationships – in other words, to determine whether and how AI directly influences children's learning or behavior. These studies allow researchers to isolate specific variables and assess their effects through controlled comparisons. For instance, one common approach is to test the "added value" of AI by comparing outcomes between children who engage with an AI-enhanced activity and those who complete the same activity without AI. This helps determine whether AI contributes something meaningful beyond the baseline experience. Other experiments may compare children's interactions with AI to those with human experts, such as one-on-one tutors. In this case, the human condition serves as a gold standard, allowing researchers to evaluate how closely AI can replicate the effectiveness of expert instruction. Other

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

studies focus on specific design features of AI – such as tone of voice, expressiveness, or responsiveness – by manipulating just one feature while holding everything else constant. This allows for a precise examination of which characteristics of AI influence child outcomes and in what ways. In all cases, choosing what to compare is not a trivial decision. It reflects different underlying research questions and requires careful methodological consideration to ensure the results are both valid and meaningful.

2.3 Study Population

Much of the current research on AI's role in child development has been conducted in regions that are relatively at the forefront of AI innovation and adoption, such as the United States, Europe, and parts of Asia. This geographic concentration might be a matter of convenience – reflecting where researchers and technical infrastructure are located – as if the study population is agnostic to the outcomes. Yet, the underlying assumption of this approach is that AI technologies are culturally neutral. In reality, AI systems are shaped by the data on which they are trained, which often encode the dominant norms, values, and biases of the societies that produce them. As a result, children do not encounter AI as a culturally blank slate but as systems embedded with implicit assumptions about language, behavior, and identity.

Two key issues highlight why AI cannot be assumed to be neutral and why study populations matter. First, AI systems may perform differently depending on a child's background. For instance, automatic speech recognition (ASR) systems have been shown to work significantly better for monolingual English-speaking children than for their bilingual peers. A recent study by Thomas et al. (2023) found that performance was lowest for bilingual children who were more dominant in their home language. Such disparities introduce inequities in children's ability to benefit from AI-enhanced educational tools. Second, beyond performance disparities, there are challenges of cultural alignment and representation. Children may treat AI agents not just as tools but as social partners – entities that speak to, respond to, and even form relationships with them. When these AI agents primarily reflect Western, mainstream cultural norms, children from marginalized or minoritized communities may find it harder to relate, resulting in reduced engagement and less effective learning. In one exemplar study, Finkelstein et al. (2013) found that children who interacted with an AI tutor that spoke African American Vernacular English (AAVE) built stronger rapport and engaged more deeply, leading to better educational outcomes.

These concerns do not arise in isolation or solely from AI technologies themselves. Rather, they reflect broader societal inequities that are encoded into

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

training data and algorithmic design. Just as teachers must navigate the complexities of supporting students from diverse backgrounds, designers and researchers must attend to how children's identities and lived experiences shape – and are shaped by – their interactions with AI. Importantly, AI may amplify these challenges in unique ways. For instance, if children view AI as an omniscient source of knowledge, they might be less likely to question the biased information generated by AI than they would with information provided by people. While many studies have acknowledged the ethical and equity implications of AI biases and misalignments, little research has formally explored how these factors might attenuate or modulate the effects AI can have on children from diverse backgrounds. This consideration is essential for framing the interpretation of the research findings discussed in the following sections, as these findings may not fully apply across all cultural, socioeconomic, or contextual settings.

3 Children's Interactions with AI Agents

3.1 How Do Children Talk to AI?

Children today interact with AI agents that vary widely in design and capability. Many of these agents can understand spoken or text-based input and generate responses, making them appear capable of conversation on the surface. Researchers studying children's development are deeply interested in how children actually engage with these agents in everyday contexts. Observing how children talk with AI offers insights not only into their usage patterns but also their interpretations and expectations of these systems.

For example, Lovato et al. (2019) recorded children's interactions with Google Assistant and analyzed the types of questions asked. They found that children aged five and six often used these agents to seek factual information – questions about science, technology, math, or practical concerns like the weather. This pattern of fact-seeking may be expected and mirrors the way many adults use voice assistants. However, Lovato et al. also documented a considerable number of socially oriented questions, in which children treated the agent as if it had personal attributes, asking things like "How old are you?" or "What's your favorite color?" A similar trend emerged in a later study on chatbots powered by generative AI with slightly older children aged eight to ten: Although most questions were factual, some focused on the AI tool's personal characteristics (Oh et al., 2025).

Children's socially framed questions – such as asking an AI tool about its favorite color or whether it has a family – have been interpreted in different ways. One common interpretation is that children are anthropomorphizing the AI tool, treating it as a social being with thoughts, preferences, or relationships. This view assumes that if children fully understood AI as a mechanical or

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

nonhuman entity, they would not ask such questions. However, this explanation may overlook a second possibility: that children engage in such questioning not only out of belief in the AI's human likeness but as a form of playful exploration. Even when aware that the AI tool is not a person, children may ask socially oriented questions to test the system's boundaries, to elicit surprising responses, or simply to enjoy the novelty of the interaction. In this view, such behaviors are driven more by curiosity and experimentation than by mistaken beliefs about the nature of AI.

Tentative support for this perspective comes from a longitudinal study. In a two-month deployment of a social robot, Kanda et al. (2007) observed that elementary-aged children initially asked many personal questions (e.g., about the robot's feelings or background), but the frequency of such questions declined over time. A more recent study involving generative AI showed a similar pattern: Children initially posed socially framed questions like "Do you go to school?" but gradually shifted toward more transactional queries, such as requests for help or information (Oh et al., 2025). If children truly believed the AI tool was a social being, one might expect such personal questions to increase over time as rapport builds and social connection deepens. Instead, the observed decline in social questions suggests that children's early interactions may often reflect playfulness and novelty seeking.

3.2 Do Children Talk to AI Differently Than They Talk to Humans?

The studies in the previous section revealed interesting patterns in children's interactions with AI agents, yet it would be challenging to draw conclusions about these behavioral patterns without directly comparing them to children's interactions with humans. To address this, researchers have employed experimental procedures where children are instructed to interact either with a person, typically an experimenter, or an AI agent. This approach allows researchers to observe children's communication from different perspectives. We will examine three aspects here: the quantity of communication, the quality of children's responses, and behaviors that reflect children's social intentions. Table 1 summarizes all the studies being discussed. It is important to note, however, that how children interact with "humans" may vary considerably depending on the context and the nature of the relationship – for example, interactions with parents may differ markedly from those with strangers. Consequently, conclusions drawn from observed AI–human differences should be approached with caution, as such differences may reflect not only the nature of AI agents versus human partners but also the relational and contextual factors shaping children's interactions in each case.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Table 1 Studies comparing children's communication with humans versus AI.

Study	Child age	N	Comparison condition	Interaction context	Communication metrics	Comparative results
Aeschlimann et al. (2020)	5–6 years	72	Voice assistant (Wizard of Oz) vs. experimenter	Game	Number of questions asked	Human = AI
					Likelihood of information sharing	Human > AI
					Likelihood of helping	Human > AI
Tolksdorf et al. (2021)	4–5 years	20	Social robot vs. experimenter	Novel word learning	Social referencing occurrences	AI > human
Xu et al. (2021)	3–6 years	90	Conversational agent vs. experimenter	Reading	Length of utterances in words	Human > AI
					Lexical diversity	Human > AI
					Topical relevance	Human > AI
					Accuracy	Human = AI
					Intelligibility	AI > human
Gampe et al. (2023)	5–6 years	50	Voice assistant (Wizard of Oz) vs. experimenter	Game	Number of questions asked	Human = AI
					Number of grounding utterances	Human > AI

					Number of syllables	Human > AI
					Pitch	Human > AI
					Volume	Human > AI
					Response latency to resolve misunderstandings	AI > human
Cohn et al. (2024)	7–12 years	71	Smart speaker vs. experimenter	Conversation	Utterance duration	AI > human
					Pitch	AI > human

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

3.2.1 Communication Quantity

The quantity of children's communication with a conversational partner refers to how frequently children participate and how much they talk. Research consistently shows that young children participate more when they interact with a human partner than with AI. For example, Xu et al. (2021) found that children aged three to six responded more often and with longer utterances during storybook reading sessions with a human experimenter than with a smart speaker. Similarly, Gampe et al. (2023), whose study involved children aged five to six interacting with either a voice assistant or a human during a treasure hunt game, observed that children used more grounding utterances when engaging with humans. These behaviors suggest higher engagement in human interactions.

3.2.2 Communication Quality

Researchers conceptualize communication quality in children's speech through multiple dimensions, focusing on both the form and function of language use such as how well children express themselves, stay on topic, and use elaborate language. These dimensions can vary depending on who the conversational partner is. On the one hand, research shows that children generally produce more sophisticated and topically relevant language when interacting with humans than with AI. For instance, Xu et al. (2021) found that children's responses to human partners were not only longer but also demonstrated greater lexical richness and stronger relevance to the topic being discussed. This suggests that children are more expressive and better able to maintain the topic of conversation when speaking with humans, perhaps because they perceive humans as partners who warrant greater communicative effort.

Other research highlights a different manifestation of communicative effort – adapting speech to facilitate the listener's comprehension. For example, in a study with children aged nine to twelve, Cohn et al. (2024) found that participants produced longer utterances and used a higher pitch when conversing with a smart speaker (Alexa) compared to a human experimenter. Such patterns suggest that children may adjust their speech to be more pronounced when interacting with AI, possibly due to being aware of its limitations as a listener and making a deliberate attempt to enhance clarity. Similarly, another study reported that children's speech is often more intelligible – clearer in articulation – when directed to a conversational agent (Xu et al., 2021). In both cases, this type of effort is expressed through adjustments to delivery to accommodate the perceived needs or limitations of AI agents.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

3.2.3 Social Intention

Researchers have also examined communication behaviors that may reflect children's social orientation, expectations, or intentions toward their communicative partner. Evidence suggests that while children can engage with AI, their interactions with human partners are marked by a deeper sense of social attunement and responsiveness. For instance, Aeschlimann et al. (2020) found that children were more likely to share information and offer help to human partners than to AI agents, suggesting that they recognize and respond to the social needs of people more than machines. Similarly, Tolksdorf et al. (2021) observed that children engaged in more frequent social referencing – looking to caregivers for guidance – when interacting with a social robot. This behavior may reflect a degree of uncertainty or hesitation in navigating interactions with AI. Adding to this, Gampe et al. (2023) showed that children were more proactive in establishing common ground with humans, using strategies such as grounding utterances and clarification efforts more often than they did with AI. Together, these findings suggest that although children interact readily with AI, their social behaviors remain more richly expressed and intentional in human communication contexts.

3.3 How Do Children React When AI Misunderstands Them?

An interesting scenario to explore in children's communication with AI is when they encounter conversational breakdowns – situations where one party misunderstands the other, leading to a deviation from the intended interaction. During child–AI interactions, these breakdowns can stem from AI's failure to comprehend a child's input or, even when it does accurately comprehend the input, it fails to respond in a way that makes sense to the conversation. However, it is important to note that conversational breakdowns are not unique to interactions between children and AI; they also frequently occur in children's communications with others. Nevertheless, it is reasonable to anticipate that such breakdowns may be exacerbated by the technical limitations of AI systems

In response to these breakdowns, children often use various strategies to restore the interaction, which include repeating their comments, adjusting their volume, articulating their words more clearly, rephrasing their statements, and even changing the subject of their inquiry (Beneteau et al., 2019; Cheng et al., 2018; Mavrina et al., 2022). For instance, Mavrina et al. (2022) found that children aged six to twelve might repeat a command if Alexa fails to understand it initially, such as repeatedly asking "Alexa, what is the most agile animal in the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

world?" after Alexa mistakenly provided information about the most venomous animal instead. Children might also adjust their volume, speaking louder to ensure the command is heard, or articulate their words more clearly, emphasizing each syllable to aid recognition. Rephrasing the statement is another common strategy, as observed when children change their request from "Alexa, what is the most agile animal in the entire world?" to "Alexa, who is the most agile?" to enhance clarity. Furthermore, if none of these strategies seem to work, children may seek help from others. For example, when Alexa failed to respond helpfully to a child's vague request – "Alexa, what can I do together with my mom and sister?" – an adult intervened by reformulating the query into a more specific question: "Alexa, events in [city name] today."

However, these studies did not compare children's repair strategies used with AI to those used with humans. Consequently, it remains unclear to what extent these strategies differ. Two experiments have provided evidence suggesting such differences are likely. The first study, conducted by Gampe et al. (2023), involved five to six-year-olds interacting with either a person or a voice assistant, both of whom exhibited staged errors in understanding children's speech. The researchers found that after a stage error occurred, children interacting with the voice assistant were less likely to continue engaging in the conversation, whereas those in the human group were more "forgiving" of their partner's errors. Thus, communication breakdowns had a more negative impact on engagement with voice assistants than with human partners.

Another study by Li et al. (2024) investigated how children aged four to eight use different repair strategies after communication breakdowns with either a human or an AI partner. Unlike the Gampe et al. (2023) study, where errors were staged, this study employed a generative AI–powered storytelling agent that responded to children in real time. This setup allowed communication breakdowns to emerge naturally during the interaction. The study confirmed that children were more likely to encounter communication breakdowns when interacting with AI than with a human partner. Additionally, children were less likely to attempt repairs following these breakdowns, even after the researchers took into consideration the children's age and language proficiency. The study specifically analyzed the strategies children used after breakdowns occurred. A notable finding was that a significant proportion of children did not resolve breakdowns caused by AI. Specifically, when the AI tool misunderstood them, they often followed the AI tool's responses rather than attempting to correct the misunderstanding. This raises questions about the developmental value of child–AI interactions, particularly in fostering social communication competence: By forgoing repair strategies – such as repetition, rephrasing, or clarification – children might miss important opportunities to practice and improve important conversation skills.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Li et al. (2024) further explored why children showed different repair behaviors depending on their conversational partner. They speculated that children are more likely to try to repair communication breakdowns when they perceive the partner as an "in-group" member – someone they feel is similar to themselves. This idea is grounded in prior research on children's motivation to engage in conversations with socially similar others (Julien et al., 2019; Sierksma et al., 2019). Indeed, the researchers found that the more children perceived AI as similar to themselves – measured by a perceived homophily scale (Rubin et al., 2020) – the more likely they were to attempt to repair communication breakdowns. These findings point to a broader issue of social alignment and motivation in children's interactions with others, which will be discussed in the following section.

3.4 Why Do Children Talk "Differently" with AI?

So far, the evidence has suggested disparities between child–AI and child–human communication, but what factors contribute to these differences? One recurring theme emerging from previous studies is that children hold differing expectations of AI's conversational capabilities, generally perceiving AI agents to be less competent in conversations than human interlocutors. Studies have found that children tend to be more articulate (Xu et al., 2021) or speak in a higher pitch (Cohn et al., 2024) when conversing with AI – likely influenced by their beliefs or prior experiences regarding AI's less-than-ideal accuracy in interpreting speech. Interestingly, children were also found to be very adept in adjusting their communications in real time based on the agent's reactions. At least one study has shown that when children interacted with either a responsive agent or one providing generic feedback regardless of what they said, their initial levels of engagement (e.g., response rate) were similar. However, as the interactions progressed, the children increasingly chose to engage with the agent that was responsive to them (Xu et al., 2024). This phenomenon aligns with theories of children's "adaptation" based on the status of their interaction partner. The ability to appropriately reciprocate or adjust to a partner's communicative response is a crucial component of communicative competence. Previous research suggests that as children grow older, their speech increasingly reflects adaptation to their partner's speech styles (Gampe et al., 2019; Street & Cappella, 1989). Even children as young as three years old have been consistently shown to possess the ability to adapt to their partner's communicative cues. It appears that such adaptations can also be transferred into the context of interactions with AI agents.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Another factor contributing to children's less motivated conversation with AI could be their perception of AI's social presence, which will be discussed greater detail in Section 4. Children may view AI as less socially present or emotionally engaging than a human, leading them to adopt more functional or goal-oriented communication (Garg & Sengupta, 2020). For example, they might use direct commands instead of engaging in conversational small talk (Kim et al., 2022). Nonetheless, It is important to note that socially oriented conversations should not necessarily be viewed as the ideal outcome for child–AI interactions. One could reasonably argue that it is beneficial for children to differentiate their modes of interaction with AI, as this capacity may help preserve a clear social boundary between AI and human partners. However, in certain contexts – such as using AI to support children's social-emotional learning or for therapeutic purposes – a lack of motivation to engage socially with AI may limit its effectiveness. In response to this challenge, developers and researchers have focused on enhancing the empathetic and affective components of AI agents. Features such as emotional validation and appraisal are being designed to encourage children to perceive AI agents not merely as tools for completing tasks but as more socially aware and emotionally responsive partners.

Children's prior experiences with technology may also shape how they respond to AI. This influence, however, appears to be determined not solely by individual variations in access and experience but also by the presence of AI in society. In particular, the cultural and technological environment in which children are embedded may shape not only their knowledge of AI but also their expectations regarding its capabilities and their attitudes toward it. In a society saturated with AI, even children with limited personal experience may develop strong expectations based on how AI is portrayed in the media, discussed in schools, or used in their communities. A particularly interesting study compared children from Pakistan (a developing country) and the Netherlands (a developed country) in their perceptions of playing with a robotic friend versus a peer (Shahid et al., 2014). The researchers found that Pakistani children tended to hold higher expectations of the robot's social capabilities and expressed greater disappointment when those expectations were unmet, for example, when the robot failed to understand child-directed speech. At the same time, they also showed greater motivation and engagement, as reflected in more animated facial expressions during play. These findings suggest that in contexts where AI is less widespread but more novel, children may approach it with both idealized expectations and heightened enthusiasm. This supports the view that societal-level factors may play an important role in shaping children's perceptions of, and interactions with, AI.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Overall, the studies reviewed employed quantitative metrics that consistently indicate lower levels of socially oriented behaviors during interactions with AI compared with human partners. Nevertheless, these differences should not be interpreted as indicating a complete absence of such behaviors in AI interactions. Rather, some children still exhibit socially oriented responses when engaging with AI, though to a lesser extent. This phenomenon was first uncovered through Nass et al.'s (1994) studies on "computers as social actors," which aimed to explain why individuals often respond socially to AI. Although their study was conducted long before AI reached its current level of sophistication, the computers they used (i.e., NeXT workstations) exhibited basic features of modern AI, albeit in rudimentary form, such as speech, language use, and responsiveness. The researchers designed a series of five studies in which participants interacted with two computers: one served as a "tutor," providing answers, while the other acted as an "evaluator," assessing the accuracy of the tutor's responses. Through these triadic interactions, the researchers found that participants applied "social rules" during the exchanges. For instance, participants tended to offer stronger praise when the computer directly solicited their feedback and exhibited gender stereotypes, perceiving praise from a male (voice) as more convincing than praise from a female (voice). These findings suggest that subtle human-like cues, particularly speech and reciprocity, prompt individuals to perceive a sense of humanity in technological systems. As a result, people unconsciously apply mental frameworks from human-to-human interactions to their engagements with technology. The findings from adult participants were later confirmed in studies with children. In one earlier study using a Wizard of Oz design (Ryokai et al. 2003), five-year-olds were assigned to take turns co-creating stories with a virtual peer named Sam or a co-present playmate. The researchers analyzed the conversations children had with Sam and found that children not only viewed Sam as a "partner" (indicating "your turn" by saying that it was Sam's turn) but also made eye contact with Sam when asking questions, much like they did with their co-present playmate. Moreover, children showed proactivity by offering help to Sam, a behavior consistent with how they interacted with human peers. These instances suggest that social elements can emerge in children's engagement with AI, and that the concept of "computers as social actors" is not necessarily incompatible with the observed differences between child–human and child–AI interactions.

3.5 Can AI Influence Children's Interactions with Others?

One reason we care about how children interact with AI is not only concern about their engagement with AI itself, but, perhaps more importantly, how these

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

interactions may shape the ways they engage with other people. In particular, since children approach AI in different ways, it is important to consider whether these patterns will carry over to their interactions with humans, potentially leading to behaviors that are perceived as unnatural or socially inappropriate. An often-discussed concern arises from the idea that when children become accustomed to giving commands to AI agents, they may carry these patterns into their interactions with people, potentially neglecting norms of politeness (Arora & Arora, 2022). There are numerous anecdotes highlighting this issue: In one blog post on *Medium*, a parent of a four-year-old, who was enthusiastic about using Alexa for things like knock-knock jokes, began to question the ramifications for the child's social etiquette. The parent noted that AI, programmed to tolerate poor manners, inadvertently encouraged the child to "boss around" (Walk, 2016). This anecdote, among many others, reflects a broader public sentiment: These commercially available AI systems are programmed to serve, with no mechanisms to hold users accountable for their verbal behavior.

Evidence beyond speculations and anecdotes to support these concerns is scarce. It takes time for children to develop norms of social interaction and etiquette through repeated interactional experiences, which makes these behaviors difficult to observe. To overcome this, some researchers have focused on a more immediate aspect: whether children can learn new linguistic routines from AI, which allows researchers to observe learning over a shorter period. In a study, children aged five to ten were instructed to talk to an AI agent designed to slow down its speech (Hiniker et al., 2021), but they could ask the AI to speed up by saying the word "Bungo." After the session, the researchers secretly instructed the children's parents to slow down their speech. Interestingly, the researchers observed that at least half of the children used the word "Bungo" to ask their parents to speed up while still in the lab. Furthermore, when the participants returned home, 55 percent of the parents reported that their child continued to use this routine, though this should be taken with a grain of salt due to the potential unreliability of self-reporting. This finding provides tentative evidence that children can pick up linguistic cues from AI. However, the results are not entirely conclusive. The word "Bungo" used in this study is relatively nonsensical and novel, which raises the question of whether children used it purely for playfulness or if it indicates a meaningful change in their linguistic patterns.

While the findings are still quite tentative, this study did point out the possibility that children could pick up linguistic routines from interacting with AI. In response to public concerns, some voice assistants, such as Amazon Alexa, have introduced features that praise children for using polite language or flag the use of less courteous expressions. Two studies offer tentative evidence

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

regarding how these features might influence children's communication behaviors. The first study involved college students who were divided into two groups: one interacted with a digital assistant that rebuked impolite requests (i.e., refusing to provide responses), while the other interacted with a control AI agent that tolerated impolite requests (Bonfert et al., 2018).[1] The researchers found that the AI agent rejecting impolite requests significantly increased users' tendency to use polite language during interactions. However, it remains unclear whether users' behavior was driven by a genuine belief that they should demonstrate politeness with AI, or if it was more motivated by the fact that they realized they would not get the information they needed without using polite language. The second study employed a different strategy (Mandagere, 2020). Instead of penalizing the non-use of polite language, they focused on positive reinforcement to praise children's use of polite language. This study observed five families and found that parents reported an increased use of polite language by children between the ages of five and thirteen. compared to the pre-test. However, there are alternative explanations that could undermine the findings due to the study's design. One issue is that the outcomes relied on parent reports, which are subject to social desirability bias or imperfect recall. Additionally, there was no comparison group, making it difficult to determine whether the observed effects were due to the encouragement from the voice assistants or other factors (e.g., simply the parents' awareness that they are part of a study focusing on politeness might have influenced their responses).

Beyond the inconclusive evidence, one potential downstream consequence of enforcing politeness in interactions with AI agents is that such features may blur the distinction between human-to-human and human-to-AI conversations, making it harder for children to differentiate between the two contexts. In addition, much public discussion has focused narrowly on politeness, such as saying "thank you" or "please", even though these expressions represent only one dimension of the broader social norms children must develop when engaging with others. Other important aspects include turn-taking, recognizing and responding to emotions, adapting communication styles to different contexts, and understanding others' perspectives. To date, we have limited evidence on how interacting with AI might alter, or perhaps not affect, such behaviors. These are behavioral routines that require time to establish, are less malleable, and often demand long-term studies to observe meaningful changes.

[1] It should be noted that this study labeled "impolite requests" as those that do not include words like "please" or do not start with common politeness markers such as "Would you" or "Could you." This itself is a quite stringent definition of politeness in the context of interaction with digital assistants.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

3.6 How Do Children Interact with AI When Others Are Present?

Most studies on children's interactions with AI have focused on one-on-one settings, including those previously discussed (e.g., Gampe et al., 2023; Xu et al., 2021). In these studies, children were typically invited into a laboratory environment to engage in structured activities such as playing a game or listening to a story with an AI agent.

However, when others, such as family members or peers, are present, children's interactions with AI can take on new dynamics. For example, consider a scenario where a smart speaker is placed in a household. The questions children ask might be overheard by parents or siblings, who could subsequently join the conversation. In addition, children may share smart toys or devices with peers at school, which makes the interactions more collaborative or negotiated. Indeed, a study focusing on families of children aged three to eight revealed that parent-only, child-only, and co-use are all common ways families engage with shared voice assistant devices at home (Wald et al., 2023). These findings suggest that children's interactions with AI are deeply embedded within the broader social and cultural contexts of their daily lives.

When parents and children are co-using AI agents such as voice assistants, these agents may play the role of mediators for managing everyday routines or regulating children's behavior, though they may also occasionally introduce tensions. Researchers tackling this question often adopt a perspective of viewing AI as an active participant within the dynamics of interaction. This approach highlights how AI, alongside humans, shapes and is shaped by the social, cultural, and technological networks it inhabits. As an example, Beneteau et al. (2020) studied how the presence of the AI voice assistant Alexa might influence how parents interact with their children in a household context. This study included children across a wide age range, from one to thirteen years old. Overall, the authors found that despite occasional conflicts caused by the smart speaker, families were able to leverage Alexa to further their parenting goals. Regarding these conflicts, the researchers documented how family members might disrupt each other's interactions with Alexa, likely for two reasons. First, Alexa was often used to facilitate tasks that did not align with another family member's goals, such as playing music the other family member disliked. This led to tensions in the shared environment where family members co-reside. Second, the sharing nature of the device, where it can only perform tasks requested by one person at a time, contributed to these disruptions. However, neither of these issues seemed to escalate beyond minor glitches. Regarding how Alexa can further parenting goals, it was used as a neutral third-party mediator for managing behaviors, such as setting timers to end an activity that

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

might otherwise have been difficult for the child to terminate. It was particularly interesting that children seemed to be more receptive to advice coming from Alexa than from their parents, likely due to the perceived objectivity resulting from its machine nature. Overall, it appears that while voice assistants might not have been designed to support family dynamics, they have, to some extent, fulfilled these unintended goals.

When peers are present, AI can also influence the nature of children's interactions, either supporting social play or distracting from it. Pretend play is a fundamental part of how young children interact with their peers, whether they are playing house, acting out adventures, or building imaginary worlds together. When AI technologies like conversational agents enter these play spaces, we may wonder: Can they become engaging playmates or will they pull children away from playing together? Pantoja et al. (2019) looked at how voice agents embodied in animal characters could support this kind of social play among three to four-year-olds. When researchers maintained control over what the agents said by typing in responses in real time, the agents could step in at just the right moments with suggestions and encouragement, which sustained children's engagement and guided them back to playing together whenever they started to drift apart. This mimics how adults such as preschool teachers or parents would naturally facilitate children's play, offering gentle prompts and support without taking over the play itself. Although this study used the Wizard of Oz technique where researchers controlled the agents behind the scenes, it still suggests that voice assistants with sophisticated contextual awareness could potentially foster meaningful social play experiences. However, when the researchers tried letting children control what the agents said themselves through a tablet app (where the children could pick topics, feelings, and events for the agent to talk about), the dynamic shifted dramatically. Children became so focused on making the agent talk – doing this almost twice every minute – that they spent less time actually playing with each other compared to when researchers controlled the agents' speech. While children showed enthusiasm for incorporating the voice agents into their play by making them talk, the tablet interface meant to give children control ended up competing for their attention rather than supporting their social interactions.

Researchers are also curious if AI agents can be designed intentionally to promote interpersonal interactions. Shared reading, for example, is a powerful way for parents and children to bond through stories, conversations, and shared experiences. Yet many educational technologies today focus on one-on-one interactions between a child and device, potentially displacing valuable parent–child interactions rather than supporting them. Xu et al. (2023) developed

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

a bilingual conversational agent, Rosita, embedded in an e-book designed to promote parent–child co-engagement in reading for families with children aged from three to six (for a demo, see https://youtu.be/UQw9j14e3yA). In addition to questions directed at children on story comprehension, Rosita asked open-ended "family questions" that invited parents and children to connect the story to their own experiences. Their studies found that family-oriented prompts encouraged substantive conversations between parents and children, encompassing daily experiences, shared memories, and discussion of the content at hand (He et al., 2025; Xu et al., 2023). This research demonstrates that when AI agents are thoughtfully designed to consider family dynamics, they can enhance rather than replace interpersonal interactions in families, turning technology into a bridge rather than a barrier to family engagement.

4 Children's Perceptions of AI

4.1 Do Children Believe AI Is a Living Being Like Humans?

When discussing AI and children, a common question people often ask is whether children believe AI is alive or real, like a person. In fact, children have long been curious about whether other things, such as plants, household machines, toys, or even fictional characters, are alive or real. This line of inquiry is grounded in their conceptual development of animate and inanimate distinction, which centers on their ontological understanding of purposes, functions, and the categorization of interactive technologies among the wide range of entities in the world (Rakison & Poulin-Dubois, 2001). A prominent theory that explains this development is Piaget's stage theory, which suggests that children progress through distinct stages of conceptual development. Initially, they may confuse animate and inanimate entities, later developing flawed distinctions (e.g., associating movement with animacy), and ultimately achieving an adult-like understanding (Opfer & Gelman, 2011). Although Piaget's theory remains influential, it has faced criticism, particularly over whether such developments occur as qualitative shifts at fixed stages or along more continuous developmental trajectories (Gelman, 1990). Despite these critiques, both Piagetian theory and perspectives advocating for continuous development agree that childhood is a period of ongoing conceptual growth, where categorizing entities like AI presents challenges. Children often focus on observable traits such as movement, associating these with "animacy" or being "alive" while overlooking the unobservable yet fundamental characteristics that may indicate AI's nonliving nature (for a review, see Opfer & Gelman, 2011).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Researchers have approached this inquiry by posing a simple question to children: What is AI? The key takeaway from this line of exploration is that children have mixed beliefs about AI falling along the animate–inanimate spectrum. In one of those studies, Xu and Warschauer (2020) asked children aged three to six years to describe a Google Assistant they had interacted with. Seventeen percent of the participants referred to it as a "human" or "girl," indicating a human-like perception, while the majority saw it as a technological object, describing it as a "device," "machine," or "phone." However, the study also found that a considerable portion of children who failed to identify the voice assistant as fitting neatly into either category, with some referring to it as "like a human but not a human" or "magic." The study also uncovered a tentative developmental difference, with all six-year-olds viewing AI firmly as an artifact, whereas those who perceived it as a living being were exclusively three-year-olds. Figure 2 shows examples of children's drawings as well as their interpretations during the drawing.

Another study focusing on slightly older children (aged six to ten years) found that, when asked what Alexa is (similar to Google Assistant), many used terms that reflected an artifact-oriented view, such as "computer chips" and "hard disks" (Festerling & Siraj, 2020). Some responses also revealed how these children positioned AI relative to humans, reasoning that humans must "put" their own intelligence into the machines they build or program, and therefore the intelligence of the "made" can never exceed that of the "maker."

Beyond asking children to categorize AI, another approach examined which perceived characteristics lead them to view it as human-like or distinctly nonhuman. These characteristics are often grouped into cognitive (thoughts), psychological (feelings or emotions), and behavioral (actions or speech) properties (Melson et al., 2009). When an entity displays either all or none of these properties, children may find it less challenging to categorize the object as either animate or inanimate. However, entities displaying only some of these properties are more likely to prompt uncertainty. Viewed as a continuum, the animate–inanimate distinction places some objects clearly at either end, while others occupy a more ambiguous middle ground.

Studies indicate that children vary in the extent to which they attribute cognitive, psychological, and behavioral capabilities to AI. For instance, Xu and Warschauer (2020) found that while most children aged three to six believed that AI has cognitive capabilities (being able to think and remember things), they were less likely to believe that AI has the capability to feel things like a friend or have emotions. When asked about their reasoning behind this, children frequently pointed to the AI agent's capacity to engage in dialogue, especially in a contingent way. In short, it appears that AI's adaptability and

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

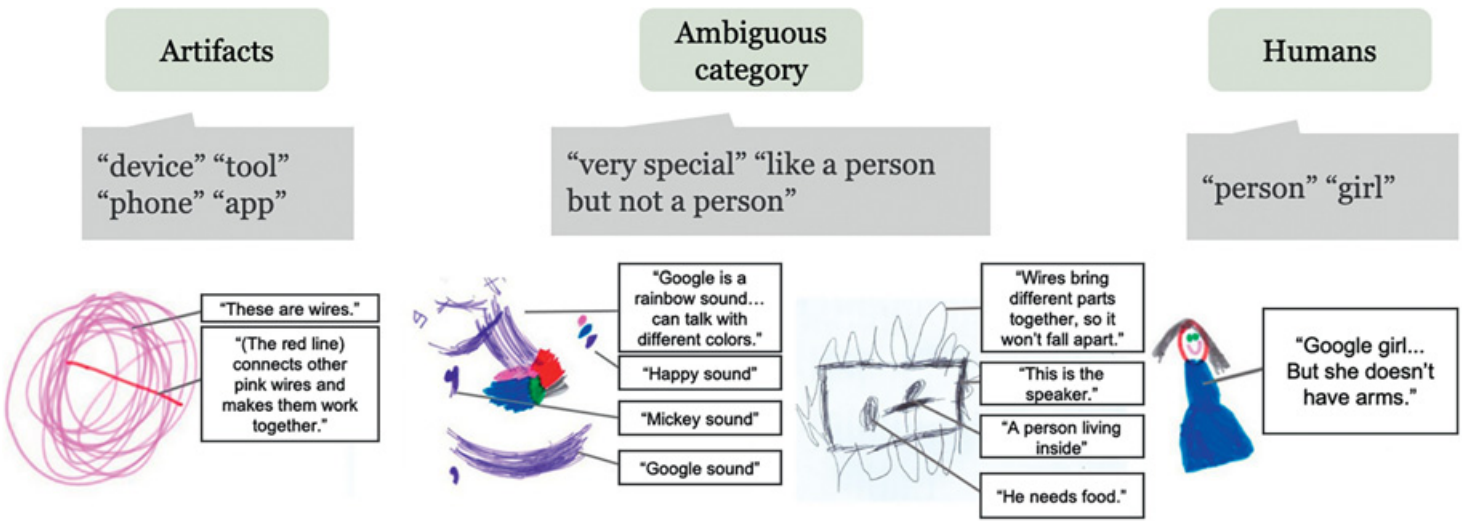


Figure 2 Children’s drawing in response to the prompt “What is inside the Google Home smart speaker?”(Xu & Warschauer, 2020).

speech capability were seen and perceived as strong indicators of intelligence, potentially contributing to the children's categorization of AI as human-like. Another study focused on children aged six to ten years and similarly found that they assigned some degree of social and mental capacity to voice assistants, and this tendency was more prominent among younger children (Girouard-Hallam et al., 2021). Interestingly, the tendency to attribute human-like characteristics to AI appears to extend beyond childhood. While few studies have explicitly compared children and adults, Cohn et al. (2024) conducted a study in which both children (aged seven to twelve) and adults interacted with Amazon Alexa and were asked how much they thought Alexa was like a real person. Across both adult and child groups, roughly 30–40 percent of participants responded affirmatively. Although the proportion was slightly higher among children, the difference did not appear to be statistically significant.

Such human-like perceptions of AI arise not only from conversational cues but also from other factors, such as physical appearance and movement. For instance, Melson et al. (2009) found that the majority of children aged seven to fifteen years affirmed that AIBO, the robotic dog, had mental states, social awareness, and moral standing. Similarly, Beran et al. (2011) suggested that a significant proportion of children between the ages of five to sixteen years in their study ascribed cognitive, behavioral, and psychological characteristics to robots. Beran et al.'s study noted that children's assigning animacy to robots is driven more by robots' physical movements than by their intelligence (though it is likely that the robot's rather repetitive patterns at that time led children to not perceive it as highly intelligent, this may not be the case with current AI). When children were asked why they considered the robot to be a living being, most children pointed to the robot's humanoid appearance and its seeming ability to move spontaneously.

Together, these findings suggest that children's perceptions of AI as human-like are shaped by a constellation of factors, including appearance (especially humanoid features), physical movement, and the capacity for contingent interaction. Rather than treating these dimensions as competing or mutually exclusive, it may be more accurate to view them as complementary cues that children weigh, often simultaneously, when making judgments about the "humanness" of AI.

4.2 Does AI Make Decisions on Its Own?

The previous section examined children's beliefs about various characteristics of AI. These observable capabilities, however, raise a deeper question: Do children perceive such behaviors as driven by programming or by an AI's

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

own "mind"? This question relates to whether children see AI as possessing a mind. People intuitively conceive of minds in terms of two broad capacities: agency and experience (Gray & Gray, 2007). *Agency* involves the ability to form intentions, reason, pursue goals, plan, communicate, and act. *Experience* encompasses the capacity to feel emotions, sense pleasure or pain, perceive through the senses, remember experiences, express a personality, and possess consciousness.

A handful of studies have used this mind perception framework to examine children's perceptions of AI agents, including voice assistants and social robots. Brink (2018) recruited children aged three to seventeen years to watch short videos of a humanoid robot and then asked them questions about their perceptions of the robot's agency and capacity for experience. Through cluster analysis, Brink suggested that over half of the child participants considered the robots to have low agency and limited capacity for feelings or emotions (experience), while another quarter viewed the robots as having high agency but still lacking capacity for experience. A more recent study used a similar method where children aged four to eleven watched videos of a mix of robots and voice assistants (Flanagan et al., 2023). The children generally ascribed low-level experiences to these agents while still attributing a certain degree of agency to them (an average of one on a scale from zero to three).

A small body of studies have compared children's mind perception of AI agents with humans as a benchmark. For instance, Flanagan et al. (2021) found that children aged five to seven, in general, attributed a similar level of ability to choose to both a robot and a human child. Another study investigated whether children understand that different people can hold individual beliefs, in contrast to conversational agents linked to the internet, which share a common set of beliefs due to synchronization (Dietz et al., 2023). Children watched videos in which a human shared different pieces of information with two separate agents – either two people or two smart speaker AI tools. The key difference was that in the human condition, each person heard different information, while in the AI condition, both devices – being connected to the internet – were assumed to share the same information. Children were then asked to judge what each agent (human or AI) would know. The findings revealed that while adults correctly inferred that AI devices would hold the same information, some children treated the AI agents more like humans, attributing individual knowledge to each. This suggests that children may have difficulty distinguishing between the belief systems of humans and AI agents.

Taken together, the studies suggest that children, to some extent, perceive AI agents as having agency, which might have implications for how they interpret their interactions with AI. For example, children might believe that AI agents

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

willingly choose to engage with them. In our own studies, many children agreed with the statement "the agent interacts with me because it chooses to do so." This belief may increase children's engagement, as the interactions feel more genuine and meaningful. However, attributing agency to AI also raises critical developmental concerns. Children's belief in AI's autonomous decision-making may lead to misconceptions about the nature of these systems, which are fundamentally governed by algorithms rather than conscious intent. Such misunderstandings can affect children's interpretation of responsibility, accountability, and the distinction between sentient beings and programmed machines. For example, when AI systems err, children might mistakenly hold the AI personally accountable rather than recognizing the limitations inherent in its design or data.

4.3 Does AI Deserve to Be Treated Fairly?

Although research in this area is still in its early stages, one emerging trend from these studies is that children generally believe it is wrong to treat AI unfairly or badly, yet they do not believe that AI deserves liberal or civil rights. One study explored this idea by creating scenarios where a humanoid robot seemed to be treated unfairly. In the study, children aged nine, twelve, and fifteen took turns playing a game with the robot (Khan Jr et al., 2012). Midway through the game, before Robovie could take its turn, an experimenter abruptly picked the robot up and shoved it into a closet, despite Robovie protesting, "It is unfair, Robovie needs to get my turn." In a post-interview, children who witnessed the disruption and the robot's protest expressed that the treatment of the robot was unjust and that the robot deserves fair treatment. However, a much smaller portion of the children considered Robovie as having civil liberties (e.g., can be sold or owned) or civil rights (e.g., can vote or be paid for work). Part of these findings were corroborated by a newer study that explicitly focused on children's considerations of fair treatment with familiar AI agents including Google Assistant and a NAO humanoid robot. On average, these children believed that it is not okay to yell at or hit an AI agent (Flanagan et al., 2023).

This pattern – children extending courtesy and care toward an embodied agent while still viewing it as object – was further supported by Newhart et al. (2016). In their study, children with chronic illnesses used robots to attend school remotely. Classmates quickly anthropomorphized the robots, referring to them by name, greeting them in the hallway, and even hugging them. Yet this human-like treatment had clear bounds. When one student, Samuel, attended picture day in person, his classmates questioned why both he and the robot were in the class photo, insisting "that's one person." Once

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Samuel was physically present, the robot was no longer seen as a legitimate stand-in and was wheeled back to its charging station. This shift underscores how children can toggle between treating the robot as socially real and categorizing it as an object, depending on context and physical presence. Indeed, children often described the robot as a teammate or friend during games, yet still talk about "parking," "recharging," or "upgrading" it when play is over (Ahumada-Newhart & Eccles, 2020; Ahumada-Newhart et al., 2023). Together, these studies reveal that children demonstrate empathy and moral reasoning in their interactions with AI, yet they very often conceptualize these agents as tools – owned and controlled rather than autonomous beings with rights.

What drives children to decide whether AI deserves moral treatment? Some might argue that the findings from previous studies – where children indicated it is not okay to treat AI agents poorly – can be seen as evidence of children demonstrating empathy toward AI, reflecting social connection. However, a recent study offers alternative explanations by exploring children's reasoning behind their moral treatment of AI in more detail (Oh et al., 2025). In this study, researchers used classic moral treatment questions, asking children whether it is okay to put Alexa in the cold or throw it in a box. After receiving their responses, they followed up with "Why?" By analyzing children's answers to these follow-up questions, the researchers found that children's concerns were less about whether AI might suffer psychological harm or be deprived of rights through unfair treatment. Instead, their reasoning often stemmed from the belief that AI's usefulness might be compromised if it is not kept in good condition, which aligns with their view of AI as a tool. For example, a child reasoned why they should treat an Alexa device gently: "Alexa wouldn't be as useful and then you are wasting plenty of money. It's like buying an entire house and smashing it." And another child similarly commented that "It might run out of battery and what are you going to do? Nothing. They are quite expensive." Thus, at first glance, children's reluctance to harm AI agents might seem like a sign of empathy toward the machines. However, a deeper analysis suggests that their behavior could instead reflect an effort to preserve the AI agent's value as a useful but fragile tool.

5 Children's Trust in AI-Generated Information

5.1 AI versus Humans: Who to Trust?

While researchers are still working to understand children's decision-making in trust toward AI, existing research on human interactions suggests a consensus that children do not blindly trust all the information they receive

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

(Harris & Corriveau, 2011). As AI becomes an increasingly integral part of how children acquire knowledge, a critical question emerges: When comparing human informants to AI informants, how do children decide whom to trust? This question falls within the broader domain of selective trust research, which examines how children evaluate and interact with different sources of information. Within this framework, two measures are commonly used to capture children's trust in human versus AI informants: endorsement and preference for seeking new information. Endorsement assesses which source – AI or human – a child is more likely to accept when the two provide conflicting information. Preference for seeking new information examines whether children choose to approach the AI or the human when they require additional information.

Stower et al. (2024) conducted a study where both a person and a robot were asked to label a novel object. Both informants were deemed reliable, as they had previously provided accurate labels for familiar objects. However, when the person and the robot gave conflicting labels for the novel object, children aged three to six years were more likely to trust the robot's label. Another study expanded on these findings by adding two additional factors: first, how children's trust in humans versus AI evolves with age, and second, whether this trust depends on the type of information being sought (Girouard-Hallam & Danovitch, 2022). Using a similar research paradigm, the study confirmed that both factors significantly influence children's trust. Specifically, children aged four to five and seven to eight were more likely to trust AI when seeking factual information (such as an animal's dietary habits). In contrast, children were more inclined to endorse and seek information from humans when the questions were personal – pertaining to the child's own information or information about other people. This pattern contextualizes Stower et al.'s (2024) results: object labeling falls within the domain of factual information, which may explain the greater trust placed in the robot in that study.

Girouard-Hallam and Danovitch (2022) also observed that the tendency to differentiate trust by information type became more pronounced with age. To explain this developmental trend, they investigated why children change the kinds of information they seek from different informants. Their findings indicated that these information-seeking behaviors and preferences were less related to children's beliefs about each informant's *general* capabilities (e.g., whether the informant can learn new things) and more closely tied to their beliefs about the informant's access to information in specific domains. Regarding AI agents, children were aware of the voice assistant's ability to access information via the internet and its broad scope of knowledge, making it better suited for factual information. In contrast, the human informant was

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

viewed as having a privileged status when it came to providing personal information about the experimenter.

Related to the discussion of trust in humans versus AI, it is also important to note that people's trust in AI is often intertwined with their trust in the creators of AI: humans. However, it is not always clear in previous research whether the intent was to capture children's trust in the technology itself or, by proxy, their trust in the humans and/or organizations perceived as responsible for implementing the technology. Research concerning adults has increasingly called for a clearer distinction between these two forms of trust (e.g., Lalot & Bertram, 2025), yet it could be quite challenging to investigate this among young children. First, many young children may not be aware that AI is created by humans in the first place, though this could be attenuated as children grow older. Indeed, a recent study on Google suggested that children at nine and ten years of age recognized that while Google serves as an effective mechanism for information gathering, the sources of information originate from humans (Girouard-Hallam & Danovitch, 2024). Second, as discussed earlier, evidence suggests that people often anthropomorphize new technologies and may therefore perceive them as direct objects of trust. Disentangling these two dimensions of trust would be a valuable direction for future research. Such work could involve in-depth interviews or the development of survey instruments designed to differentiate between trust in technology and trust in its creators, as well as to explore the interplay between these forms of trust.

5.2 What Factors Influence Children's Trust in AI ?

In addition to understanding how children make decisions between AI and human informants, it is also important to acknowledge that variations within AI informants, such as differences in how they deliver information and how they present their identity, can influence the extent to which children choose to trust the information provided.

5.2.1 History of Accuracy in the Past

A key factor that influences children's trust in AI-generated information is whether it has been able to provide accurate information in the past. Just like children's trust in human informants (see Harris & Corriveau, 2011), children are also more likely to trust a robot who has a consistent track record of being accurate over one that does not. For instance, to examine the role of past accuracy in young children's trust of robots, Brink and Wellman (2020) first established a reliability contrast by having one robot consistently label familiar objects correctly and the other incorrectly. Three-year-old participants

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

were then shown a novel object and asked which robot they would rather ask for its name. Each robot then provided a different name for the novel object, and the children were asked to indicate which name they believed was correct. They found that three-year-olds trusted information from an accurate robot over an inaccurate one. Reliance on an AI agent's history of accuracy as a cue to reliability can help children obtain trustworthy information and reduce the risk of being misled, particularly when they lack sufficient prior knowledge or relevant experience. However, an overreliance on past accuracy may also overshadow other relevant cues, leading children to trust an informant solely because of their prior reliability, even when the current claim may be inaccurate.

Given that past accuracy strongly shapes children's trust, an important question is what happens when a previously reliable AI agent makes mistakes or provides inaccurate information. Will children then lose trust in it? Evidence suggests that children's trust does decline when such agents make mistakes (Weiss et al., 2010), but it can be restored if the AI agent subsequently demonstrates consistent accuracy. Di Dio et al. (2020) investigated how children's trust toward a human and a robot changed across three phases: trust acquisition, trust loss, and trust restoration. Children aged three to nine were asked to judge whether their play partner correctly guessed which box contained a hidden doll. In the trust acquisition phase, researchers measured how many trials it took for children to begin consistently following their partner's guesses, defining trust as doing so for three consecutive trials. The trust loss phase measured how many incorrect trials were needed before children stopped following their partner's guesses. The trust restoration phase replicated the procedure of the initial trust acquisition phase to assess whether and how quickly trust could be rebuilt. The results showed that children's trust in the robot could be restored, but it required more trials than were needed to establish trust initially, a pattern that was also observed with human partners.

Other factors can nonetheless interfere with children's ability to rely on past accuracy when judging the trustworthiness of AI agents. For example, as discussed earlier, Brink and Wellman (2020) also found that three-year-olds were more likely to trust an anthropomorphic robot with a history of accuracy than one without such a record. However, this effect disappeared if the robot lacked human-like features: children showed no difference in trust regardless of the robot's past accuracy. This contrast suggests that three-year-olds may only use past accuracy as a trust cue when the AI agent has human-like attributes, indicating that these attributes can carry more weight than performance history in their trust decisions. Similarly, Baumann et al. (2023) compared children's selective trust in a competent robot and a human who had

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

provided inaccurate inforamtion. Three-year-olds trusted the incompetent human significantly more than the robot, whereas five-year-olds trusted the competent robot more. This shift suggests that, with age, children become increasingly able to prioritize epistemic cues over conflicting social or perceptual cues when evaluating whom to trust.

5.2.2 Familiarity and Attachment

When children seek information from people, their choices are influenced not only by reliability in the past but also by other social factors. Researchers speculate that these social factors may similarly shape children's selective trust in AI agents. One factor is *familiarity*, which might influence how receptive children are to nonhuman informants. While there have been no studies explicitly examining the role of familiarity with AI agents, research shows that children tend to trust familiar media characters more than unfamiliar ones. For example, in one study, four-year-old children were presented with factual statements about novel animals, plants, foods, and activities by either a familiar character (e.g., Nemo from *Finding Nemo*) or an unfamiliar counterpart resembling the familiar character (Danovitch & Mills, 2014). Children generally preferred to seek answers from the familiar character and endorsed its statements over those of the unfamiliar character, even when the latter was similar in appearance. Although children's trust in the familiar character diminished somewhat when it provided inaccurate information, their preference did not fully shift to the accurate but unfamiliar character.

At a deeper level, children's preference for familiar characters may not stem solely from familiarity itself but could be influenced by parasocial relationships – one-sided emotional bonds formed through repeated interactions (Calvert, 2017). When deciding whom to trust, children may engage both analytic and heuristic modes of processing. In this context, strong emotional cues, such as familiarity, can activate heuristic processing, in which affect plays a central role in guiding decisions. In the case of familiarity and attachment to AI agents, these mechanisms are likely to influence children's interactions in similar ways – potentially even more strongly. This is likely due to AI agents' ability to reciprocally respond to children (which we will discuss next), which could potentially foster a stronger sense of bonding.

5.2.3 Communicative Behaviors

Children's trust in AI agents is also shaped by the communicative cues embedded in the interaction. For example, one communicative behavior

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

studied involved an AI agent engaging in small talk, a socially conventional practice that signals an intention to maintain reciprocal communication and can, in turn, facilitate the development of trusting relationships. Van Straten et al. (2022) examined this by comparing children's overall trust in a humanoid NAO robot (instead of their trust in the specific information it provided, as was investigated in the previously mentioned studies). Children aged seven to ten interacted with robots that either did or did not engage in small talk (e.g., "How are you?" or "How old are you?"). The results showed that when NAO asked such questions, the children reported higher overall trust in the robot. Further analysis suggested that this increase in trust likely stemmed from the children's heightened perception of the robot's ability to align with their perspective.

AI agents' confidence level in delivering responses also influences children's trust, and is typically communicated through a number of linguistic and paralinguistic cues, including tone, posture, and word choice. For instance, in a study focusing on human informants, simply adding "I think" to testimony caused three-year-olds to become skeptical of the provided responses (Jaswal & Malone, 2007). Another study corroborated this and found that using certain language (e.g., "I know … ") significantly increased three- to six-year-old children's trust in the speaker, compared to statements using uncertain language (e.g., "I guess … ", Kertesz et al., 2021). It is possible that the confidence level expressed by AI agents might similarly influence children's evaluations. Evidence from one study with adults supports this expectation: when an AI system prefaced its response with an expression of uncertainty (e.g., "I am not sure, but"), participants reported lower confidence in the system and were less likely to agree with its answers. At the same time, these expressions of uncertainty increased participants' accuracy in identifying incorrect responses. AI agents such as ChatGPT often deliver information with a high degree of confidence – both in tone and in the way they present their knowledge – even when the answer may be inaccurate or unknowable.[2] This tendency could inadvertently inflate users' trust in the system, potentially leading them to accept its answers without sufficient scrutiny.

6 AI and Children's Learning

In the previous section, we discussed children's interactions with AI, as well as their perceptions and trust in these agents. These moment-to-moment interactions

[2] More recent research (Steyvers et al., 2025) shows that newer large language models are improving: They can better distinguish between answerable and unanswerable questions, and also recognize the boundaries of their own knowledge.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

have important implications for learning – particularly for how children acquire skills or knowledge, transfer them to other domains, and apply them to everyday situations over time. In what follows, we shift focus to the domain of learning itself, examining how these interactions may serve not only as social exchanges but also as potential foundations – or limitations – for children's learning from AI.

6.1 Can AI Help Children Learn?

One area people often look at when exploring AI's role in learning is how it can support children's language learning. Children develop language through interactive conversations with others – whether parents, teachers, or peers (see Rowe & Snow, 2020). As AI systems become more capable of supporting natural communication, they offer new opportunities for language exposure and practice. AI researchers have developed systems that function as language partners, simulating human-like interactions in activities like reading, storytelling, and role-playing. For example, Xu et al. (2022a) developed a conversational agent that carried out dialogic reading with children aged three to six, narrating stories, asking questions, and providing feedback based on their responses. When comparing story comprehension outcomes, the study found that children who read with the AI agent showed similar improvements to those who read with human partners. Beyond verbal interactions, AI can also simulate nonverbal cues – such as eye gaze and gestures – that support language development. Westlund et al. (2017) compared how preschoolers responded to directional cues from both a robot and a human when learning the names of unfamiliar animals. The study revealed that children were equally adept at interpreting both the robot's and the human's eye gaze and postural signals to identify referenced objects, achieving comparable word retention rates.

More recent studies have relied on generative AI to support more free-form language activities, such as co-creating stories. Fan et al. (2024) developed an interactive visual storytelling platform where children from second to fourth grade could customize main characters by defining their traits and emotions and select story scenes. The system then used AI-generated keywords to guide children in constructing story plots for each scene. Children used the system to create stories of different genres, including fantasy, science fiction, and adventure. They rated both the system and the quality of the stories highly, with the AI-generated images being their favorite feature. This study demonstrated how generative AI has the potential to scaffold children's creative writing process while maintaining their agency and enthusiasm through multimodal interactions that combine visual and narrative elements.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

AI-assisted language learning has also been studied among children with, or at risk of, language difficulties, as well as among multilingual learners. Estévez et al. (2021) evaluated a NAO robot's effectiveness in speech therapy sessions over thirty weeks with five children aged nine to twelve who had language disorders to improve their reading comprehension, pronunciation, phonological awareness, and other foundational literacy skills. While therapists were present in the room, the children interacted directly with the NAO robot, which engaged them in a wide range of activities such as story retelling, answering comprehension questions, and identifying syllable structures by providing structural guidance, support, and positive reinforcement. Therapists observed children's improvement in areas including vocalization, memorization, and sentence construction through sustained attention and motivation to engage in the activities with NAO. For multilingual learners, AI could help address challenges arising from limited exposure to one of their languages at home – particularly when parents are not fully bilingual – or from insufficient opportunities to practice both languages in educational settings that prioritize one language over another. Xu et al. (2023) designed a bilingual conversational agent named Rosita that could engage in shared reading of culturally responsive books in either Spanish or English while comprehending input in both languages. This flexibility allowed three to six-year-olds to naturally code-switch between languages during conversations, supporting the development of two languages simultaneously in the context relevant to them. In a similar vein, Lee and Jeon (2022) designed a conversational agent for slightly older English language learners (aged seven to nine) to engage in small talk, object guessing games, and preposition learning activities through structured dialogue in which the agent posed questions and provided responsive feedback. Most children perceived this agent as a human-like partner.

While AI shows promise in supporting children's language development, challenges remain in replicating crucial aspects of human–child interaction. Joint attention – the shared focus between child and communicative partner on objects or events – is fundamental to early language learning (Tomasello & Farrar, 1986). During shared reading, for example, this occurs naturally when a parent notices their child's interest, such as gazing at a picture, and builds upon it through questions and engagement. Parents instinctively respond to subtle behavioral cues such as a child leaning forward in interest or turning away in disengagement, but creating AI systems that can recognize and appropriately respond to these dynamic attention signals remains technically challenging. One study attempted to address this challenge through a system that tracks joint attention between parent and child using computer vision for gaze detection and speech recognition to identify verbal references to objects

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

(Kwon et al., 2022). It then recommended topics on a screen to promote contextually relevant conversation between the parent and the child aged one to three. While the system increased parent utterances responsive to children's attention by 38.3 percent compared to static guidance cards, the system only worked on the predefined objects and required three cameras, making it hard to scale.

Establishing common ground presents another hurdle, as effective language development relies on building shared understanding through personal connections (Bohn & Köymen, 2018). While parents naturally reference a child's experiences – such as relating a story about animals to the child's recent zoo visit – AI systems struggle to maintain this rich contextual awareness. Though AI can be programmed with general knowledge about child development (e.g., Dietz Smith et al., 2024; Zhang et al., 2024), it cannot easily replicate how human caregivers adjust their interactions based on intimate knowledge of a child's world. Advancing these capabilities will require improvements in both natural language processing and the ability to gather and utilize personal context while maintaining privacy and security.

A summary of the reviewed studies is presented in Table 2.

Besides the studies in Table 2 that focused on language learning outcomes, a growing number of studies have explored how children can learn other subject-matter knowledge – such as math, science, and computational thinking – through similar mechanisms, particularly dialogue-based interactions with AI designed to support learning. While we do not aim to provide an exhaustive list here, we highlight a few examples to illustrate how this kind of learning can extend to other subject areas. For example, Xu et al. (2022) integrated conversational agents into children's television shows produced by PBS KIDS, allowing students aged three to six to engage in dialogue with characters about science concepts during the narrative, leading to better science learning compared to passive viewing. Similarly, Dietz et al. (2021) developed StoryCoder, a voice-guided mobile application that teaches computational thinking to children aged five to eight through interactive storytelling activities. Their evaluation with twenty-two children showed that participants could successfully learn computational concepts such as sequences and loops through story customization games. Recent advances in generative AI has enabled more open-ended learning through storytelling, as it can engage in dynamic, free-form conversations that adapt to children's responses while maintaining educational objectives through proper prompt engineering. The Mathemyths system, developed by Zhang et al. (2024), provides an example of this capability. Using GPT-4, it co-creates stories with children while naturally weaving mathematical

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Table 2 Example studies on AI as children's language learning partner.

Study	Application	Study design	N	Findings
Xu et al. (2022b)	A conversational agent for dialogic reading	2 (dialogic reading or nondialogic reading) × 2 (agent or human) factorial design	117 children aged 3−6	A conversational agent can replicate the benefits of dialogic reading with a human partner.
Westlund et al. (2017)	A social robot for word learning	Within-subject experiment (learning with a robot vs. a human)	36 children aged 2−5	Children recalled new words equally well from a robot versus from a human.
Fan et al. (2024)	An interactive system that allows children to co-create stories with generative AI	User study	8 children aged 6−10	Children successfully used the system to create stories of different genres (e.g., sci-fi, adventure, fantasy) and perceived the system very well.
Estévez et al. (2021)	A NAO robot that engages children with language difficulties in language-related activities	Longitudinal study	5 children aged 9−12	Therapists observed children's improvement in vocalization, memorization, and sentence construction through sustained attention and motivation to engage in the activities with NAO.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Table 2 (cont.)

Study	Application	Study design	N	Findings
Xu et al. (2023)	A bilingual conversational agent for parent–child shared reading	User study	18 parent–child dyads with children aged 3−6	The bilingual agent effectively engaged children verbally and encouraged parental involvement while supporting their linguistic flexibility.
Lee & Jeon (2022)	A conversational agent for English language learners	User study	76 children aged 9	The majority of the students considered the agent as having human characteristics and perceived it as a human-like language partner.
Kwon et al. (2022)	A system that tracks visual and spoken language cues of parents and children to infer joint attention and recommend topics for discussion	Within-subject experiment (having the system vs. static paper guidance cards)	8 parent–child dyads with children aged 1−3	Parent verbal responses to their child's focus of attention improved by 38.3 percent when using the system versus traditional guidance cards.

vocabulary into the narrative. Unlike earlier systems that relied on pre-scripted responses, Mathemyths can generate contextually relevant continuations of children's story ideas while integrating target mathematical vocabulary, providing a more personalized and creative learning experience. Children achieved comparable math learning outcomes through interactions with Mathemyths versus a human.

6.1.1 Key Design Features in AI Systems for Learning

Across domains – from language development and literacy to subject-matter learning – AI systems that support children's learning often share core design features contribute to their effectiveness. Despite differences in purpose, modality, and learner age, many employ principles grounded in developmental research and pedagogical best practices. We will now highlight several features that recur across studies.

Structured scaffolding: Structured scaffolding refers to the intentional sequencing of supports – such as prompts, hints, feedback, and multimodal cues – that keep a task within a child's optimal difficulty range and gradually fade as the child gains mastery. For example, in dialogic reading and video-based learning, conversational agents adapted their questions when children struggled, rephrasing them into simpler, multiple-choice formats. This adjustment significantly improved preschoolers' story comprehension and conceptual understanding (Xu et al., 2022a, 2022b). Similarly, when children displayed misconceptions, an intelligent math tutor provided targeted hints that guided learners toward the correct reasoning without directly supplying the answer (Zhang & Chen, 2022). In both cases, carefully staged guidance helped sustain attention and break down complex learning goals into manageable steps, allowing children to build understanding incrementally.

Human involvement: Parental or peer involvement transforms the AI–child interaction into a triadic learning experience, where the human partner offers emotional connection, responsiveness, and real-world context, while the AI agent delivers sustained, adaptive practice. In a bilingual shared-reading system, for instance, parents built on the AI agent's prompts by linking the story to their child's everyday life – elaborating on key ideas and vocabulary introduced by the system (Xu et al., 2023). Similarly, a gaze-tracking prompt system enabled collaborative engagement by notifying parents when their toddler fixated on an object, prompting them to ask timely and relevant questions. This significantly increased back-and-forth, responsive conversation between parents and children (Kwon et al., 2022). Even in therapeutic and classroom settings, the triadic model proves valuable: During speech-language sessions

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

with a NAO robot, therapists were able to step back and observe the child's interaction, fine-tuning goals while the robot led exercises and gave feedback (Estévez et al., 2021). Across these examples, human–AI partnerships leverage the natural social motivation children feel toward familiar adults and peers, deepening engagement and grounding new knowledge in rich, shared experiences.

Child-driven, creative interactions: Open-ended creative interaction foregrounds children's autonomy, empowering them to originate ideas – whether narratives, visual scenes, or algorithmic solutions – and then iteratively refine them with AI support, transforming the act of *making* into a vehicle for *learning*. In Mathemyths, a GPT-4 tutor integrates mathematical terms into the narrative children create, linking new vocabulary to self-generated plot lines and thereby cementing learning (Zhang et al., 2024). Fan et al.'s multimodal platform lets second- to fourth-graders specify characters, emotions, and settings, after which generative AI produces illustrated sequences; as they critique and adjust these outputs, children practice genre structure, descriptive language, and visual storytelling (Fan et al., 2024). Similarly, in the StoryCoder system designed for boosting computational thinking, five- to eight-year-olds decide how on-screen figures should solve problems, and the agent converts their choices into executable code blocks (Dietz et al., 2021). Across such systems, the child's creative autonomy – paired with an AI agent that supplies timely prompts, feedback, and multimodal renderings – cultivates curiosity, self-efficacy, and durable transfer of knowledge and skills.

Inclusive design approaches: Effective child–AI systems recognize that learners have diverse languages, cultural identities, and developmental profiles, and design the content and interaction style accordingly. Bilingual conversational agents like Rosita demonstrate how language and cultural alignment can deepen engagement and learning. Designed to resemble the cultural backgrounds of the families it serves, Rosita reads culturally responsive books in both English and Spanish, allowing children to code-switch naturally and parents to participate in ways that feel authentic and comfortable (Xu et al., 2023). This not only reinforces children's language development in both languages but also validates their home-language practices and identities. For children with language disorders, embodied AI systems such as the NAO robot offer targeted, consistent support through structured, repetitive activities. In a thirty-week intervention, NAO guided children through tasks such as sound identification, syllable segmentation, and sentence construction, using clear prompts and immediate feedback (Estévez et al., 2021). By delivering individualized practice at a steady pace, NAO effectively supported core language skills in children with special needs. Together, these examples show how inclusive

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

design – attuned to children's linguistic, cultural, and developmental contexts – can make AI systems more responsive, equitable, and effective learning partners.

6.2 Does Using AI Impact Children's Ability to Learn?

In the preceding sections, we examined how AI can support children's language learning and potentially other subject-specific knowledge. Outcomes such as vocabulary learning or science concept understanding are typically viewed as proximal outcomes because they are directly linked to the instructional content AI delivers. In this role, AI functions as a teaching agent, providing explanations, posing questions, and offering feedback – much like a human teacher. When used in this way, AI tends to produce consistent improvements in specific learning domains. Yet, beyond these domain-specific outcomes, it is also important to examine how AI might foster the broader foundational skills that support learning and development across domains (McCoy & Sabol, 2024).

6.2.1 Curiosity

Broadly defined as a desire to learn or explore the unknown, curiosity plays a critical role in children's learning. It is often manifested through exploratory behaviors and question-asking. Although curiosity as a stable trait can be difficult to measure, researchers typically focus on situational or "state" curiosity – the form that can be observed in specific moments of learning and shaped by the surrounding environment.

AI has the potential to both satisfy and stimulate this kind of curiosity. As a rich source of information and an always-available conversational partner, AI can respond to children's questions but also encourage them to pursue further inquiry. In a two-week study of children's interactions with a generative AI agent, Oh et al. (2025) described "passages of intellectual search" – moments when children continued to ask increasingly specific follow-up questions either because the answer left gaps that prompted further probing or because the AI' agent's response introduced new information that stimulated additional questioning. These instances demonstrate how AI can not only fulfill children's immediate informational needs but also afford deeper and more sustained exploration.

Other studies aim to stimulate epistemic curiosity through two primary approaches: by modeling curiosity-related behaviors and by highlighting the gap between a child's current knowledge and a desired understanding. These strategies are often used in combination. For example, Goren et al. (2015) designed an interaction using a Story-Maker app co-played on a tablet by the child and a social robot. In this setup, the child manipulated characters on the screen, while

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

the robot narrated a story that matched the child's actions. One group of children interacted with a *curious robot* – a robot that expressed enthusiasm for learning and exploration, challenged the child, and suggested novel actions in the app. Another group interacted with a *neutral robot* that behaved as a cooperative playmate but did not display overt curiosity or exploration-related behaviors. The study involved forty-eight children between the ages of three and eight. The results showed that children who played with the curious robot exhibited significantly more free exploration and uncertainty-seeking behaviors compared to those who played with the neutral robot. These findings support the idea that curiosity can be socially transmitted – that is, observing curiosity-related behaviors can evoke similar behaviors in children. This result was corroborated by another study focused on slightly older elementary school children interacting with a conversational agent (Abdelghani et al., 2022). Together, these findings suggest that modeling curiosity through AI can be an effective strategy for creating curiosity-inducing contexts that promote exploratory and information-seeking behaviors in children. However, it remains unclear how much exposure – or what "dosage" or intensity – to such interactions would be necessary for curiosity-promoting experiences to evolve into stable dispositions in children.

6.2.2 Creativity

There is a lack of consensus on how to conceptualize and measure creativity as an ability, and this disagreement contributes to the difficulty of providing definitive answers to the question of whether AI impacts creativity. The most comprehensive empirical evidence so far comes from a study in which adults were asked to engage in a series of creativity problem-solving tasks, such as coming up with novel uses for everyday items. For example, one task asked participants to use paper clips, water bottles, and paper bags to create a new toy. Their ideas were then rated based on originality and appropriateness – how practical the idea was. The researchers found that participants who used AI were able to generate ideas that were rated as both more original and appropriate than those who used a Google web search (Lee & Chung, 2024). Yet such enhancement in individual-level creativity appears to come at the expense of diversity at the collective level (Doshi & Hauser, 2024). That is, while individuals may generate ideas that are novel in isolation, these ideas collectively tend to exhibit greater homogeneity, raising concerns about the long-term implications for creativity at the societal level, particularly as generative AI becomes a common anchor for ideation.

For studies involving children, most of them designed specialized child-facing interfaces for AI agents with the goal of boosting creativity, particularly

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

creative expressions like storytelling and drawing. One principle utilized in these interfaces was to use AI agents to *model* good creative behaviors, which children could potentially internalize as they observed and engaged in interactions with the AI partner. Following this principle, Elgarf et al. (2021) developed a creative storyteller agent with an accompanying screen-based interface. The agent was designed to embody key qualities of creative expression, such as generating ideas fluently and sharing surprising twists in the plot lines. The researchers then compared the story ideas generated by seven- to eight-year-olds interacting with this agent to those of children interacting with another agent that did not model these creative behaviors. When comparing the average creativity scores of the two groups, the researchers did not find evidence that the creative agent group performed better. While this is subject to a host of alternative explanations such as how the creativity was measured and how the statistical analysis was done, it is possible that behavioral modeling alone, without explicit guidance, may not be sufficient to influence children's behaviors.

As a result, other studies have adopted a more direct approach, in which the agent's role is to provide explicit scaffolding, which potentially helps make the typically abstract and fluid creative process more concrete and structured. In a system developed by Zhang et al. (2022), children aged six to ten were asked to generate creative stories, with the AI agent generating drawings to visualize the children's narration as an external representation. The system helps both to concretize and visualize children's ideas, as well as to provide additional visualizations that support subsequent idea generation. For the former, the AI system automatically recognizes the entities mentioned by children in their narration. For example, when a child says "two pink flowers on the roadside," the system extracts "two," "pink," and "flower" to create sketches in real time. For the latter, the system builds upon the existing ideas and generates additional sketches designed to spark imagination and encourage the development of further ideas based on the figurative drawings. The researchers found that adding this support component significantly boosted children's fluent generation of story ideas, as well as the level of elaboration in those ideas. There are a number of similar programs that have quite consistently demonstrated that providing this kind of explicit scaffolding boosts children's performance in creativity tasks.

However, the question remains whether AI itself possesses creativity. Developmental psychology researchers have compared the "creativity" demonstrated by children to that generated by AI and argued that human children are much better at coming up with unpredictable yet feasible ideas. To explain this difference, researchers like Gopnik (2022) have attributed AI's inherent

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

limitations to the fact that, as children grow, they build abstract models of the world around them. These abstract frameworks enable children to make predictions and generalizations that significantly diverge from existing information, leading to more dramatic departures from prior knowledge – something fundamentally different from how AI predicts the next token in a sequence. This raises the following question: If AI lacks this capability, is it convincing to many that AI can genuinely promote creativity? One way to address this tension is by rethinking AI's role – not as a source of creativity, but as a facilitator that establishes conditions under which children's creativity can flourish. In this sense, AI-based creativity support tools do not act as creative agents in their own; rather, they reduce cognitive load and lower barriers to ancillary tasks in creative work. By doing so, they can amplify children's creative performance, even if the creativity itself originates in the child.

6.2.3 Self-regulation

When discussing AI's impact on children's regulation processes, two interrelated constructs are worth clarifying. The first is self-regulation, which is closely related to executive functioning. It refers to the capacity to monitor and manage internal states, emotions, and behaviors in the service of achieving a particular goal. It includes components such as inhibitory control, working memory, and cognitive flexibility. In the meantime, self-regulated learning (SRL) refers to a learner's ability to actively manage the learning process. This includes planning, monitoring one's understanding, evaluating progress, and reflecting on strategies. While self-regulation underpins SRL, the latter is more task-specific and typically occurs within academic contexts. When it comes to AI's role, these two concepts are often discussed in different but overlapping ways. On the one hand, researchers consider how children's self-regulation abilities may shape the way they interact with AI. On the other, there is growing interest in how AI can scaffold and support SRL.

Studies exploring AI as a scaffold for SRL – particularly those involving real-time, in-context prompts – suggest that AI can be designed to help children monitor and reflect on their progress. For example, Weng et al. (2024) developed a system that integrated AI feedback to support students' SRL. Their findings show that AI can help make regulation strategies more explicit, prompting primary school students to reflect in the moment and adopt habits such as routinely asking themselves metacognitive monitoring questions. However, open questions remain about whether children are able to internalize these self-regulation strategies once the AI support is removed. This concern is particularly relevant given that children will likely encounter AI in contexts

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

where such structured assistance is absent. One promising direction is to design AI systems not only to provide immediate scaffolding but also to gradually teach children how to regulate their own learning. When structured in this way, AI can begin by modeling regulatory strategies, then encourage learners to take on increasing responsibility for monitoring and control, and ultimately fade its support as children internalize these skills. Evidence from such a prototype system in mathematics learning suggests that this gradual transfer approach can improve regulation accuracy, reduce dependence on AI prompts, and support the transfer of self-regulation strategies to new contexts (Molenaar, 2022).

7 Children's Learning about AI

7.1 What Does AI Literacy Entail?

A key characteristic of AI systems is their opacity, as children are generally unaware of what happens between their input (e.g., what they say to AI agents) and the system's output (e.g., how AI agents respond). This issue is compounded by children's still developing cognitive abilities and limited knowledge of technology. Many child-facing AI products are designed to emulate humans, encouraging children to perceive them as playmates, coaches, or companions, often without revealing their internal workings. Such black-box designs can inadvertently encourage overreliance on AI and limit opportunities for children to critically evaluate or question the information they receive. In response, the research community has increasingly focused on understanding and addressing these challenges by focusing on promoting children's literacy, that is, the skills needed to interact with AI effectively.

AI literacy is often considered a subset or extension of digital literacy, which rose to prominence in the public sphere two decades ago with the advent of the internet, smart devices, and social media (Long & Magerko, 2020; Yang, 2022). Similar to the ongoing debates surrounding digital literacy, AI literacy remains an evolving area of discussion regarding the core skills that constitute this construct. Scholars have proposed that AI literacy can be understood through a multidimensional framework that combines different ways of defining both *AI* and *literacy*. Specifically, AI can be conceptualized as technical systems (e.g., machine learning foundations), as tools (e.g., Google Assistant), and as sociotechnical systems with societal impacts (e.g., implications for children's privacy). Literacy, in turn, may be functional (e.g., understanding and implementing machine learning models), critical (e.g., critically evaluating AI products and using them responsibly), and inherently motivational (e.g., fostering interest in STEM). Based on this framework, researchers have proposed more specific skill sets for AI literacy. For example, Long and Magerko (2020) define AI literacy as

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

a set of skills that enable individuals to critically evaluate, communicate, and collaborate effectively with AI as a tool. This definition encompasses five broad areas of inquiry: What is AI? What can AI do? How does AI work? How should AI be used? How do people perceive AI? It should be noted that both of these frameworks consider "children" with an unspecified age range, likely encompassing kindergarten to high school students. Other frameworks have been developed that are more tailored to specific age groups. For example, Su et al. (2023) synthesized AI literacy for early childhood through a systematic review of sixteen published articles in this context. Their proposed scope of AI literacy for early childhood covered dimensions similar to those outlined in broader, age-agnostic frameworks by Gu and Ericson (2025) and by Long and Magerko (2020), such as basic AI concepts (e.g., machine learning, rule-based systems). While there may be differences in the depth of each competency area, the overall range of topics remains consistent across these frameworks.

Based on these conceptual frameworks, some early studies have developed assessments aimed at measuring children's AI literacy. Most of these assessments focus on evaluating children's knowledge of either the technical aspects of AI (e.g., machine learning and algorithms) or recognizing its societal implications (e.g., privacy concerns, job displacement, or fairness in decision-making). These often take the form of quizzes, surveys, or structured interviews that gauge what children know or believe about AI. However, relatively few assessments have directly evaluated children's ability to engage with AI through hands-on activities – such as interacting with AI systems, building simple models, or making informed decisions using AI tools – which may offer a more nuanced picture of AI literacy in practice. In this section, we will outline several example assessments targeting children.

Studies have found that kindergarten children can understand basic AI knowledge and concepts, such as how AI is trained by humans using data, how AI can help solve problems in daily life, how AI has certain limitations in what it can do, and knowledge of different AI algorithms (e.g., Su et al., 2023; Williams et al., 2019). While children can understand these fundamentals, there is still a lack of reliable AI literacy assessments for this age group. Most existing assessments have been tied to evaluating specific AI literacy curricula. For instance, Williams et al. (2019) developed an AI platform called PopBots where children aged four to six interacted with a robot and a tablet app to learn three AI concepts: knowledge-based systems, supervised machine learning, and generative AI. Leveraging this system, the researchers developed a set of ten multiple-choice questions that asked children to predict the robot's behavior in child-friendly contexts. For example, in one item, children are told that a robot has been taught to classify

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

strawberries and tomatoes as belonging to the "good" group, and that it has not been taught about any other foods. They are then asked which group the robot would place chocolate in: the good group or the bad group. The correct answer is "good," since the robot has only seen examples of items in the good group. This question assesses whether children understand how the algorithm is initialized, and whether they recognize that a robot exposed only to positive or negative examples will assign all new items to that one category.

Likely because AI literacy education has become more prevalent starting from elementary school, with structured curricula covering both the functional and societal aspects of AI, more comprehensive assessments have been developed and validated for this age group to measure different constructs of AI literacy. Chung et al. (2025) designed an assessment for elementary school students that covers three key areas: understanding AI concepts (e.g., "What is AI?", "What can AI do?"), how AI works (e.g., knowledge representation, decision-making, machine learning), and AI ethics. The researchers created twenty-four modular questions using real-life scenarios that elementary students would encounter, such as identifying AI applications in daily life or understanding how AI makes decisions. This assessment tool was validated through multiple rounds of expert review (n=15) and pilot testing with 287 sixth-grade students. The results showed good reliability (KR-20 score > 0.81) and appropriate difficulty levels for the target age group.

AI literacy assessments for younger students focus on foundational concepts, and those designed for middle and high school students place greater emphasis on technical terminology and more advanced understandings of the definitions, functions, and ethical implications of AI. Chiu et al. (2024) developed a twenty-five-item multiple-choice test that measures middle school students' AI literacy and validated this measure among 2,390 students in Hong Kong. This assessment aligns with the AI curriculum framework implemented in K–12 education in Hong Kong, which includes three learning areas: knowledge of AI, process in AI, and impact of AI. In the knowledge of AI dimension, students must differentiate between AI and non-AI technologies and understand core concepts such as neural networks and natural language processing. The process in AI dimension tests students' knowledge of machine learning approaches through items about reinforcement learning, model training concepts, and the importance of training data. The impact of AI dimension examines understanding of ethical and societal implications through items about job displacement, transparency, responsible development, and privacy. Table 3 presents an example item for each dimension.

While these instruments could be valuable tools for gaining an initial understanding of children's knowledge about AI, it is unclear whether they

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Table 3 Example items in the AI literacy assessment developed by Chiu et al. (2024).

Dimension	Example item
Knowledge of AI	3. Which of the following does NOT apply AI technology? A. Web browsing B. Facial recognition C. Semantic analysis D. Speech recognition
Process in AI	12. Which of the following does adopt the trial-and-error way? A. Reinforcement learning B. Supervised learning C. Unsupervised learning D. Deep learning
AI impact	23. Which of the following operations is considered appropriate when building computer vision applications? A. Collecting and storing sensitive data during video surveillance. B. Faking others' identity during facial recognition. C. Developing medical imaging applications without tests and ethical reviews. D. Making safety the priority when building auto-driving cars.

adequately capture children's ability to *use AI as tools*. These skills are likely difficult to measure through quiz-style items and may be more reliably assessed through scenario-based instruments that require children to actually engage with AI tools. In addition, because most AI literacy assessments are grounded in specific curricular frameworks, they tend to focus on a particular subset of skills. This scope makes it difficult to use these instruments to capture children's general abilities in AI literacy.

7.2 Can Young Children Learn AI Literacy and Ethics?

Many research programs have focused on developing curricula or platforms to explicitly teach children AI literacy, with the primary aim of helping children acquire technical knowledge as well as an understanding of the societal and ethical implications of AI. Many of these programs provide children with hands-on experiences, such as engaging with AI algorithms or designing AI systems for social purposes, and have shown that by participating in such

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

educational opportunities, children's knowledge and attitudes toward AI, along with their awareness of its broader societal impacts, can be significantly improved. For example, in one program, a child-friendly app was designed to allow middle-schoolers to work alongside their parents to create their own conversational agents. The app employed a building-block design, allowing children to add functions to the conversational agents by stacking blocks instead of requiring them to understand coding syntax. An evaluation after the program suggested that this learning opportunity boosted children's confidence in determining when their agents should be trusted, likely due to their heightened understanding of the inner workings of AI (Van Brummelen et al., 2023).

Another strand of work has emphasized children's ability to understand fairness and bias in AI. For instance, in a cross-cultural study, Charisi et al. (2021) explored how children in Japan and Uganda conceptualized fairness, including fairness in robot-related scenarios. Children as young as six were able to express concerns about biased behavior by robotic companions. For instance, Ugandan children emphasized fairness in physical and material terms, imagining scenarios where a robot should ensure equal access to food, clothing, or classroom entry. Japanese children, on the other hand, highlighted psychological aspects of fairness, such as a robot's ability to understand a student's reasons for being late and to judge accordingly. These examples demonstrate that even young children can apply culturally grounded ideas of fairness to their expectations of intelligent systems. Moreover, children's understanding of AI bias can be deepened through interactive tools. For example, Melsion et al. (2021) developed an interactive platform that explained how its image classification system generated predictions. By visualizing the features driving its decision-making, the system improved children's understanding of algorithmic bias and helped them recognize discriminatory patterns in AI outputs. Together, these examples suggest that AI literacy education can be approached not only through teaching technical foundations but also by fostering children's awareness of fairness, bias, and the ethical challenges posed by intelligent systems.

8 Conclusion

This Element has provided an overview of some pressing questions regarding AI's implications for children's development. While there are many pathways to approach this broad issue, we attempted to focus on these implications across three interconnected domains: how children interact, how they perceive, and how they learn. First, interaction behaviors serve as the primary way children engage with AI, allowing them to gain direct experience that shapes their perceptions. Second, perception reflects children's internal understanding of

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

AI, which is influenced by and also guides their behaviors. Third, learning emerges as the consequence of these interactions and perceptions, encompassing not only children's knowledge and skills about AI itself but also the broader developmental outcomes fostered through AI-facilitated experiences.

Across these domains, a recurring throughline is how children's engagement with AI both parallels and diverges from their interactions with human partners. On the surface, children's interactions with AI often resemble those with humans – they ask questions, seek help, and even make emotional remarks. Yet, nuances still exist, likely stemming from children's differing expectations regarding the responsiveness and trustworthiness of AI. This dilemma is also reflected in children's perceptions: Some children blur the line between humans and machines – sometimes attributing thoughts and emotions to AI – while others recognize AI as fundamentally different from humans. These perceptions, in turn, influence how children learn with AI. When children treat AI as knowledgeable but impersonal, their learning strategies may differ from those they use with teachers, peers, or parents. Some may defer uncritically to AI as an all-knowing entity; others may engage through experimentation and a small amount of trial and error. Crucially, children's learning from AI does not occur in isolation – it is shaped by how they understand the social and cognitive limits of AI compared to humans. Overall, framing concerns about AI in terms of whether it replaces humans could oversimplify the issue. Indeed, throughout our discussions, it is clear that AI can elicit many similarities with humans, but it remains fundamentally different. It provides unique sources of experience and information – distinct yet complementary to the other resources children already have.

The evidence so far suggests that while access to AI can offer significant advantages to children, it also raises important considerations for mitigating potential risks. Below are a few bullet points that have emerged from the evidence, illustrating the complex coexistence of both opportunities and risks. We will elaborate on some examples:

- **Personalized, adaptive learning at scale.** AI systems, such as conversational agents and intelligent tutoring platforms, have demonstrated the ability to deliver personalized instruction that adjusts to children's developmental levels, language backgrounds, and learning needs. For instance, dialogic reading agents improve story comprehension comparable to human partners, and STEM-focused AI supports concept understanding through interactive narratives. These systems can reach learners who might otherwise have limited access to tailored educational support.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

- **Stimulation of curiosity and creativity.** Research reveals that AI agents modeling curiosity and exploratory behavior can significantly increase children's engagement in learning activities. Creative storytelling systems scaffold idea generation and elaboration, supporting children's creative expression and cognitive growth in ways that mirror human scaffolding. At the same time, concerns remain that dependence on AI-simulated curiosity and creativity may inadvertently constrain the development of children's own agency.
- **Social engagement with others.** On the one hand, there are legitimate concerns that increased time spent with AI may displace human interaction and reduce opportunities for engagement with others. On the other hand, AI technologies, when designed for collaborative use with adults or peers, may also facilitate social interaction. For example, AI-supported shared reading encourages parent–child conversations, and voice assistants have been observed mediating family routines, supporting behavior management, and facilitating engagement without replacing interpersonal relationships.
- **Anthropomorphism and misattribution of agency.** Children's conceptualizations of AI often blur the boundaries between human and machine. Many children attribute thoughts, emotions, and intentions to AI agents, sometimes viewing them as "like a human but not a human" or as magical entities. While this can increase engagement, it may also lead to misunderstandings about AI's capabilities and limitations, potentially fostering overtrust.
- **Limitations in social and emotional engagement**. Empirical findings show that children's communication with AI is less socially rich compared to interactions with humans. Children tend to share less information, offer less help, and engage in fewer conversational repairs with AI. Moreover, after AI misunderstandings, children are less likely to attempt repairs and often adjust their communication to fit AI limitations, missing critical opportunities to practice social communication skills that are essential for development.
- **Embedded bias and cultural misalignment**. AI systems can encode societal biases and fail to perform equitably across diverse populations. For example, automatic speech recognition models have lower accuracy for bilingual children, especially those dominant in a non-English home language. Furthermore, AI agents that primarily reflect mainstream cultural norms may reduce engagement for children from marginalized communities, limiting the effectiveness of AI-based learning tools.
- **Privacy, safety, and content integrity concerns.** AI's data collection, reliance on internet connectivity, and content generation capabilities raise

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

significant concerns regarding children's privacy and exposure to inappropriate or inaccurate content. Current legal frameworks often lag behind technological advances, leaving gaps in protections specific to AI's unique risks.

DNonetheless, these guiding principles are not intended as a definitive or exhaustive list, but rather as initial examples within a rapidly evolving landscape of AI design. Furthermore, it is crucial to consider how these insights can be meaningfully translated into the design and implementation of child-facing AI systems, including decisions about whether such systems should be used by children at all. A productive way forward might be to use these overarching principles or foundational pillars while also actively engaging children, families, and community members in co-design processes. This includes conducting field tests and efficacy trials to ensure that systems not only reflect theories but also align with the lived experiences, needs, and values of those they are intended to serve. By combining both theory-driven insights and child-centered approaches, we can hopefully strike a balance between advancing AI capabilities and safeguarding children's developmental needs. This balanced approach has the potential to ensure that AI not only augments children's learning and social development but also aligns with ethical standards that prioritize their wellbeing in an increasingly AI-driven world.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

References

Abdelghani, R., Oudeyer, P. Y., Law, E., de Vulpillières, C., & Sauzéon, H. (2022). Conversational agents for fostering curiosity-driven learning in children. *International Journal of Human-Computer Studies*, *167*, 102887.

Aeschlimann, S., Bleiker, M., Wechner, M., & Gampe, A. (2020). Communicative and social consequences of interactions with voice assistants. *Computers in Human Behavior*, *112*, 106466.

Ahumada-Newhart, V., & Eccles, J. S. (2020). A theoretical and qualitative approach to evaluating children's robot-mediated levels of presence. *Technology, Mind, and Behavior*, *1*(1).

Ahumada-Newhart, V., Schneider, M., & Riek, L. D. (2023). The power of robot-mediated play: Forming friendships and expressing identity. *ACM Transactions on Human–Robot Interaction*, *12*(4), 1–21.

Arora, A., & Arora, A. (2022). Effects of smart voice control devices on children: Current challenges and future perspectives. *Archives of Disease in Childhood*, *107*(12), 1129–1130.

Baumann, A.-E., Goldman, E. J., Meltzer, A., & Poulin-Dubois, D. (2023). People do not always know best: Preschoolers' trust in social robots. *Journal of Cognition and Development*, *24*(4), 535–562.

Beneteau, E., Boone, A., Wu, Y., et al. (2020). Parenting with Alexa: Exploring the introduction of smart speakers on family dynamics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery.

Beneteau, E., Richards, O. K., Zhang, M., et al. (2019). Communication breakdowns between families and Alexa. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery.

Beran, T. N., Ramirez-Serrano, A., Kuzyk, R., Fior, M., & Nugent, S. (2011). Understanding how children understand robots: Perceived animism in child–robot interaction. *International Journal of Human-Computer Studies*, *69*(7–8), 539–550.

Bickham, D., Schwamm, S., & Izenman, E. (2024). *Use of voice assistants and generative AI by children and families*. The Digital Wellness Lab.

Bohn, M., & Köymen, B. (2018). Common ground and development. *Child Development Perspectives*, *12*(2), 104–108.

Bonfert, M., Spliethöver, M., Arzaroli, R., et al. (2018). If you ask nicely: A digital assistant rebuking impolite voice commands. In *Proceedings of*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

the 20th ACM International Conference on Multimodal Interaction (pp. 95–102). Association for Computing Machinery.

Brink, K. (2018). A theory of (the technological) mind: Developing understanding of robot minds (Doctoral dissertation, University of Michigan). http://deepblue.lib.umich.edu/handle/2027.42/146010.

Brink, K. A., & Wellman, H. M. (2020). Robot teachers for children? Young children trust robots depending on their perceived accuracy and agency. *Developmental Psychology*, *56*(7), 1268–1277.

Calvert, S. L. (2017). Parasocial relationships with media characters: Imaginary companions for young children's social and cognitive development. In F. C. Blumberg & P. J. Brooks (Eds.), *Cognitive development in digital contexts* (pp. 93–117). Elsevier Academic Press. https://doi.org/10.1016/B978-0-12-809481-5.00005-5.

Charisi, V., Imai, T., Rinta, T., Nakhayenze, J. M., & Gomez, R. (2021, June). Exploring the concept of fairness in everyday, imaginary and robot scenarios: A cross-cultural study with children in Japan and Uganda. In *Proceedings of the 20th Annual ACM Interaction Design and Children Conference* (pp. 532–536). Association for Computing Machinery.

Cheng, Y., Yen, K., Chen, Y., Chen, S., & Hiniker, A. (2018). Why doesn't it work? Voice-driven interfaces and young children's communication repair strategies. In *Proceedings of the 17th ACM Conference on Interaction Design and Children* (pp. 337–348). Association for Computing Machinery.

Chiu, T. K. F., Chen, Y., Yau, K. W., et al. (2024). Developing and validating measures for AI literacy tests: From self-reported to objective measures. *Computers and Education: Artificial Intelligence*, *7*, 100282.

Chung, K., Kim, S., Jang, Y., Choi, S., & Kim, H. (2025). Developing an AI literacy diagnostic tool for elementary school students. *Education and Information Technologies*, *30*(1), 1013–1044.

Cohn, M., Barreda, S., Graf Estes, K., Yu, Z., & Zellou, G. (2024). Children and adults produce distinct technology- and human-directed speech. *Scientific Reports*, *14*(1), 15611.

Danovitch, J. H., & Mills, C. M. (2014). How familiar characters influence children's judgments about information and products. *Journal of Experimental Child Psychology*, *128*, 1–20.

Di Dio, C., Manzi, F., Peretti, G., et al. (2020). Shall I trust you? From child–robot interaction to trusting relationships. *Frontiers in Psychology*, *11*, Article 469. https://doi.org/10.3389/fpsyg.2020.00469.

Dietz, G., Le, J. K., Tamer, N., et al. (2021). StoryCoder: Teaching computational thinking concepts through storytelling in a voice-guided app for

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

children. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). Association for Computing Machinery.

Dietz, G., Outa, J., Lowe, L., Landay, J. A., & Gweon, H. (2023). Theory of AI mind: How adults and children reason about the "mental states" of conversational AI. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *45*(45). https://escholarship.org/uc/item/8x88d8tp.

Dietz Smith, G., Prasad, S., Davidson, M. J., Findlater, L., & Shapiro, R. B. (2024). ContextQ: Generated questions to support meaningful parent-child dialogue while co-reading. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference* (pp. 408–423). Association for Computing Machinery.

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, *10*(28), eadn5290.

Egert, F., Cordes, A.-K., & Hartig, F. (2022). Can e-books foster child language? Meta-analysis on the effectiveness of e-book interventions in early childhood education and care. *Educational Research Review*, *37*, 100472.

Elgarf, M., Skantze, G., & Peters, C. (2021). Once upon a story: Can a creative storyteller robot stimulate creativity in children? In *Proceedings of the 21th ACM International Conference on Intelligent Virtual Agents* (pp. 60–67). Association for Computing Machinery.

Estévez, D., Terrón-López, M.-J., Velasco-Quintana, P. J., Rodríguez-Jiménez, R.-M., & Álvarez-Manzano, V. (2021). A case study of a robot-assisted speech therapy for children with language disorders. *Sustainability*, *13*(5), 2771.

Fan, M., Cui, X., Hao, J., et al. (2024). StoryPrompt: Exploring the design space of an AI-empowered creative storytelling system for elementary children. In *Extended abstracts of the CHI Conference on Human Factors in Computing Systems*, New York, NY, USA (pp. 1–8). Association for Computing Machinery.

Festerling, J., & Siraj, I. (2020). Alexa, what are you? Exploring primary school children's ontological perceptions of digital voice assistants in open interactions. *Human Development*, *64*(1), 26–43.

Finkelstein, S., Yarzebinski, E., Vaughn, C., Ogan, A., & Cassell, J. (2013). The effects of culturally congruent educational technologies on student achievement. In H. C. Lane, K. Yacef, J. Mostow, & P. Pavlik (Eds.), *Artificial intelligence in education* (pp. 493–502). Springer.

Flanagan, T., Rottman, J., & Howard, L. H. (2021). Constrained choice: Children's and adults' attribution of choice to a humanoid robot. *Cognitive Science*, *45*(10), e13043.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Flanagan, T., Wong, G., & Kushnir, T. (2023). The minds of machines: Children's beliefs about the experiences, thoughts, and morals of familiar interactive technologies. *Developmental Psychology*, *59*(6), 1017–1031.

Gampe, A., Zahner-Ritter, K., Müller, J. J., & Schmid, S. R. (2023). How children speak with their voice assistant Sila depends on what they think about her. *Computers in Human Behavior*, *143*, 107693.

Gampe, A., Wermelinger, S., & Daum, M. M. (2019). Bilingual children adapt to the needs of their communication partners, monolinguals do not. *Child Development*, *90*(1), 98-107.

Garg, R., & Sengupta, S. (2020). He is just like me: A study of the long-term use of smart speakers by parents and children. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, *4*(1), Article 11, 1–24.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and the animate–inanimate distinction as examples. *Cognitive Science*, *14*(1), 79–106.

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford University Press.

Girouard-Hallam, L. N., & Danovitch, J. H. (2022). Children's trust in and learning from voice assistants. *Developmental Psychology*, *58*(4), 646–661.

Girouard-Hallam, L. N., & Danovitch, J. H. (2024). How does Google get its information? Children's judgements about Google search. *British Journal of Developmental Psychology*, *42*(3), 334–347.

Girouard-Hallam, L. N., Streble, H. M., & Danovitch, J. H. (2021). Children's mental, social, and moral attributions toward a familiar digital voice assistant. *Human Behavior and Emerging Technologies*, *3*(5), 1118–1131.

Glaser, M., Garsoffky, B., & Schwan, S. (2009). Narrative-based learning: Possible benefits and problems. *Communications*, *34*(4), 429–447.

Golinkoff, R. M., Hoff, E., Rowe, M. L., Tamis-LeMonda, C. S., & Hirsh-Pasek, K. (2019). Language matters: Denying the existence of the 30-million-word gap has serious consequences. *Child Development*, *90*(3), 985–992.

Gopnik, A. (2022). Children, creativity, and the real key to intelligence. *APS Observer*, 35. www.psychologicalscience.org/observer/children-creativity-intelligence.

Gordon, G., Breazeal, C., & Engel, S. (2015, March). Can children catch curiosity from a social robot?. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human–Robot Interaction* (pp. 91–98). Association for Computing Machinery.

Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, *315*(5812), 619.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Gu, X., & Ericson, B. J. (2025). AI literacy in K-12 and higher education in the wake of generative AI: An integrative review. In *ICER '25: Proceedings of the 2025 ACM Conference on International Computing Education Research* V. 1 (pp. 125–140). Association for Computing Machinery.

Harris, P. L., & Corriveau, K. H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 1179–1187.

He, K., Cervera, K., Levine, J., et al. (2025). Promoting parent–child shared reading with a bilingual conversational agent. *International Journal of Human-Computer Studies*, *199*, 103489.

Hiniker, A., Wang, A., Tran, J., et al. (2021). Can conversational agents change the way children talk to people? In *Interaction design and children*, Athens, Greece (pp. 338–349). ACM .

Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development*, *8*(3), 263–283.

Julien, H. M., Finestack, L. H., & Reichle, J. (2019). Requests for communication repair produced by typically developing preschool-age children. *Journal of Speech, Language, and Hearing Research*, *62*(6), 1823-1838.

Kahn Jr, P. H., Kanda, T., Ishiguro, H., et al. (2012). "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. *Developmental Psychology*, *48*(2), 303–314.

Kanda, T., Sato, R., Saiwaki, N., & Ishiguro, H. (2007). A two-month field trial in an elementary school for long-term human–robot interaction. *IEEE Transactions on Robotics*, *23*(5), 962–971.

Kertesz, A. F., Alvarez, J., Afraymovich, M., & Sullivan, J. (2021). The role of accent and speaker certainty in children's selective trust. *Cognitive Development*, *60*, 101114.

Khan Academy. (2024). *Annual report 2023–2024: Khanmigo*. https://annualreport.khanacademy.org/khanmigo.

Klarin, J., Hoff, E. V., Larsson, A., & Daukantaitė, D. (2024). Adolescents' use and perceived usefulness of generative AI for schoolwork: Exploring their relationships with executive functioning and academic achievement. *Frontiers in Artificial Intelligence*, *7*, 1415782.

Kim, M. K., Druga, S., Esmaeili, S., et al. (2022). Examining voice assistants in the context of children's speech. *International Journal of Child-Computer Interaction*, *34*, 100540.

Kwon, T., Jeong, M., Ko, E.-S., & Lee, Y. (2022). Captivate! Contextual language guidance for parent–child interaction. In *Proceedings of the 2022*

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

CHI Conference on Human Factors in Computing Systems (pp. 1–17). Association for Computing Machinery.

Lalot, F., & Bertram, A.-M. (2025). When the bot walks the talk: Investigating the foundations of trust in an artificial intelligence (AI) chatbot. *Journal of Experimental Psychology: General*, *154*(2), 533–551.

Lee, B. C., & Chung, J. (2024). An empirical investigation of the impact of ChatGPT on creativity. *Nature Human Behaviour*, *8*(10), 1906–1914.

Lee, S., & Jeon, J. (2024). Visualizing a disembodied agent: Young EFL learners' perceptions of voice-controlled conversational agents as language partners. *Computer Assisted Language Learning*, *37*(5–6), 1048–1073.

Li, Z., Thomas, T., Yu, C.-L., & Xu, Y. (2024). "I said knight, not night!": Children's communication breakdowns and repairs with AI versus human partners. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference* (pp. 781–788). Association for Computing Machinery.

Ligthart, M., Van Bindsbergen, K. L. A., Fernhout, T., et al. (2019). A child and a robot getting acquainted: Interaction design for eliciting self-disclosure. In *18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2019* (pp. 61–70). International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery.

Lovato, S. B., Piper, A. M., & Wartella, E. A. (2019, June). Hey Google, do unicorns exist? Conversational agents as a path to answers to children's questions. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children* (pp. 301–313). Association for Computing Machinery

Madden, M., Calvin, A., Hasse, A., & Lenhart, A. (2024). *The dawn of the AI era: Teens, parents, and the adoption of generative AI at home and school*. Common Sense.

Mandagere, P. (2020). The effects of AI on children's use of politeness (Honors thesis). Western Michigan University.

Mavrina, L., Szczuka, J., Strathmann, C., et al. (2022). "Alexa, you're really stupid": A longitudinal field study on communication breakdowns between family members and a voice assistant. *Frontiers in Computer Science*, 4. https://doi.org/10.3389/fcomp.2022.791704.

McCoy, D. C., & Sabol, T. J. (2025). Overcoming the streetlight effect: Shining light on the foundations of learning and development in early childhood. *American Psychologist*, *80*(2), 135–147.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Melson, G. F., Kahn Jr., P. H., et al. (2009). Children's behavior toward and understanding of robotic and living dogs. *Journal of Applied Developmental Psychology*, *30*(2), 92–102.

Melsión, G. I., Torre, I., Vidal, E., & Leite, I. (2021). Using explainability to help children understand gender bias in AI. In *Interaction Design and Children*, Athens Greece (pp. 87–99). ACM.

Molenaar, I. (2022). The concept of hybrid human–AI regulation: Exemplifying how to support young learners' self-regulated learning. *Computers and Education: Artificial Intelligence*, *3*, 100070.

Munzer, T. G., Miller, A. L., Weeks, H. M., Kaciroti, N., & Radesky, J. (2019). Differences in parent–toddler interactions with electronic versus print books. *Pediatrics*, *143*(4), e20182012.

Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 72–78). Association for Computing Machinery.

Newhart, V. A., Warschauer, M., & Sender, L. (2016). Virtual inclusion via telepresence robots in the classroom: An exploratory case study. *The International Journal of Technologies in Learning*, *23*(4), 9–25.

OECD. (2019). *Education at a Glance 2019: OECD Indicators*. OECD. https://doi.org/10.1787/f8d7880d-en.

Oh, S., Zhang, C., Girouard-Hallam, L., Zhou, Z., March, H., Jayaramu, S., & Xu, Y. (2025). "Hey Curio, can you tell me more?": Children's information-seeking and trust in AI. In *Proceedings of the ACM International Conference on Interaction Design and Children (IDC '25)*. Association for Computing Machinery .

Oh, S., Zhang, T., & Xu, Y. (2025). Empathy or utility: Children's reasoning for moral consideration of AI. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*.

Opfer, J. E., & Gelman, S. A. (2011). Development of the animate–inanimate distinction. In U. Goswami (Ed.), *The Wiley-Blackwell handbook of childhood cognitive development* (pp. 213–238). John Wiley & Sons.

Pantoja, L. S., Diederich, K., Crawford, L., & Hourcade, J. P. (2019). Voice agents supporting high-quality social play. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children* (pp. 314–325). Association for Computing Machinery.

Rakison, D. H., & Poulin-Dubois, D. (2001). Developmental origin of the animate–inanimate distinction. *Psychological Bulletin*, *127*(2), 209–228.

Rowe, M. L., & Snow, C. E. (2020). Analyzing input quality along three dimensions: Interactive, linguistic, and conceptual. *Journal of Child Language*, *47*(1), 5–21.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Rowland, C. F. (2007). Explaining errors in children's questions. *Cognition*, *104*(1), 106–134.

Ryokai, K., Vaucelle, C., & Cassell, J. (2003). Virtual peers as partners in storytelling and literacy learning. *Journal of Computer Assisted Learning*, *19*(2), 195–208.

Rubin, R. B., Palmgreen, P., & Sypher, H. E. (2020). Perceived homophily measure. In R. B. Rubin, P. Palmgreen, & H. E. Sypher (Eds.), *Communication research measures* (pp. 278–281). Routledge.

Ryokai, K., Vaucelle, C., & Cassell, J. (2003). Virtual peers as partners in storytelling and literacy learning. *Journal of Computer Assisted Learning*, *19*(2), 195–208.

Sierksma, J., Spaltman, M., & Lansu, T. A. (2019). Children tell more prosocial lies in favor of in-group than out-group peers. *Developmental Psychology*, *55* (7), 1428–1439.

Shahid, S., Krahmer, E., & Swerts, M. (2014). Child–robot interaction across cultures: How does playing a game with a social robot compare to playing a game alone or with a friend?. *Computers in Human Behavior*, *40*, 86–100.

Steyvers, M., Tejeda, H., Kumar, A., et al.(2025). What large language models know and what people think they know. *Nature Machine Intelligence*, *7*, 221–231.

Street, R. L., & Cappella, J. N. (1989). Social and linguistic factors influencing adaptation in children's speech. *Journal of Psycholinguistic Research*, *18*, 497–519.

Stower, R., Kappas, A., & Sommer, K. (2024). When is it right for a robot to be wrong? Children trust a robot over a human in a selective trust task. *Computers in Human Behavior*, *157*, 108229.

Su, J., Ng, D. T. K., & Chu, S. K. W. (2023). Artificial intelligence (AI) literacy in early childhood education: The challenges and opportunities. *Computers and Education: Artificial Intelligence*, *4*, 100124.

Thomas, T., Takahesu-Tabori, A., Stoehr, A., Varady, C., & Xu, Y. (2023) Does bilingual status influence automatic speech recognition for young Latino children? Paper presented at the *14th International Symposium on Bilingualism*, Sydney, Australia, June 26–30, 2023.

Tolksdorf, N. F., Crawshaw, C. E., & Rohlfing, K. J. (2021). Comparing the effects of a different social partner (social robot vs. human) on children's social referencing in interaction. *Frontiers in Education*, *5*. https://doi.org/10.3389/feduc.2020.569615.

Tomasello, M., & Farrar, M. J. (1986). Joint attention and early language. *Child Development*, *57*(6), 1454–1463.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Tong, Y., Wang, F., & Danovitch, J. (2020). The role of epistemic and social characteristics in children's selective trust: Three meta-analyses. *Developmental Science*, *23*(2), e12895.

Van Brummelen, J., Tian, M. C., Kelleher, M., & Nguyen, N. H. (2023, September). Learning affects trust: Design recommendations and concepts for teaching children – and nearly anyone – about conversational agents. In B. Williams, Y. Chen, & J. Neville (Eds.), *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 37, No. 13, pp. 15860–15868). AAAI Press.

van Straten, C. L., Peter, J., Kühne, R., & Barco, A. (2022). On sharing and caring: Investigating the effects of a robot's self-disclosure and question-asking on children's robot perceptions and child–robot relationship formation. *Computers in Human Behavior*, *129*, 107135.

Wald, R., Piotrowski, J. T., Araujo, T., & van Oosten, J. M. F. (2023). Virtual assistants in the family home: Understanding parents' motivations to use virtual assistants with their child(dren). *Computers in Human Behavior*, *139*, 107526.

Walk, H. (20, April 4). The Amazon Echo is turning my kid into an asshole. https://onezero.medium.com/amazon-echo-is-magical-it-s-also-turning-my-kid-into-an-asshole-dc49ec66872f.

Weiss, A., Wurhofer, D., & Tscheligi, M. (2009). "I love this dog" – Children's emotional attachment to the robotic dog AIBO. *International Journal of Social Robotics*, *1*(3), 243–248.

Weng, X., Xia, Q., Ahmad, Z., & Chiu, T. K. (2024). Personality traits for self-regulated learning with generative artificial intelligence: The case of ChatGPT. *Computers and Education: Artificial Intelligence*, *7*, 100315.

Westlund, J. M. K., Dickens, L., Jeong, S., Harris, P. L., DeSteno, D., & Breazeal, C. L. (2017). Children use non-verbal cues to learn new words from robots as well as people. *International Journal of Child-Computer Interaction*, *13*, 1–9.

Williams, R., Park, H. W., Oh, L., & Breazeal, C. (2019). PopBots: Designing an artificial intelligence curriculum for early childhood education. *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(1), 9729–9736.

Xu, Y. (2023). Talking with machines: Can conversational technologies serve as children's social partners? *Child Development Perspectives*, *17*(1), 53–58.

Xu, Y., Aubele, J., Vigil, V., Bustamante, A. S., Kim, Y.-S., & Warschauer, M. (2022a). Dialogue with a conversational agent promotes children's story comprehension via enhancing engagement. *Child Development*, *93*(2), e149–e167.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Xu, Y., He, K., Levine, J., et al. (2024). Artificial intelligence enhances children's science learning from television shows. *Journal of Educational Psychology*, *116*(7), 1071–1092.

Xu, Y., He, K., Vigil, V., et al. (2023). "Rosita reads with my family": Developing a bilingual conversational agent to support parent–child shared reading. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference* (pp. 160–172). Association for Computing Machinery.

Xu, Y., Vigil, V., Bustamante, A. S., & Warschauer, M. (2022b). Contingent interaction with a television character promotes children's science learning and engagement. *Journal of Applied Developmental Psychology*, *81*, 101439.

Xu, Y., Wang, D., Collins, P., Lee, H., & Warschauer, M. (2021). Same benefits, different communication patterns: Comparing children's reading with a conversational agent vs. a human partner. *Computers & Education*, *161*, 104059.

Xu, Y., & Warschauer, M. (2020). What are you talking to? Understanding children's perceptions of conversational agents. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery.

Yang, W. (2022). Artificial intelligence education for young children: Why, what, and how in curriculum design and implementation. *CAEAI*, *3*, 1–7.

Zhang, C., Liu, X., Ziska, K., et al.(2024). Mathemyths: Leveraging large language models to teach mathematical language through child–AI co-creative storytelling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1–23). Association for Computing Machinery.

Zhang, S., & Chen, X. (2022). Applying artificial intelligence into early childhood math education: Lesson design and course effect. In *2022 IEEE International Conference on Teaching*, *Assessment and Learning for Engineering (TALE)* (pp. 635–638). Institute of Electrical and Electronics Engineers (IEEE).

Zhang, C., Yao, C., Wu, J. et al. (2022). StoryDrawer: A child–AI collaborative drawing system to support children's creative visual storytelling. In *CHI Conference on Human Factors in Computing Systems*, New Orleans, LA, USA (pp. 1–15). Association for Computing Machinery.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Acknowledgements

This research is supported by the Overdeck Family Foundation and the National Science Foundation under Grant #2115382. We thank Harvard Graduate School of Education Monroe C. Gutman Library for supporting the Gold Open Access fee for this publication.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Cambridge Elements

Generative AI in Education

Mark Warschauer

University of California, Irvine

Mark Warschauer is Distinguished Professor of Education at the University of California, Irvine, with affiliated faculty appointments in the Departments of Informatics, Language Science, and Psychological Science. He is a member of the National Academy of Education and the director of the UCI Digital Learning Lab. Professor Warschauer is one of the most influential scholars in the world on digital learning, digital literacy, and the use of AI in education. He has published 12 books on these topics including with MIT Press, Cambridge University Press, Teachers College Press, and Yale University Press, and some 300 scientific articles and papers. His work has been cited more than 48,000 times, making him one of the most cited researchers in the world on educational technology. He previously served as founding editor of *Language Learning & Technology* and inaugural editor of *AERA Open*.

Tamara Tate

University of California, Irvine

Tamara Tate is a Project Scientist at the University of California, Irvine, and Associate Director of the Digital Learning Lab. She leads the Lab's work on digital and online tools to support teaching and learning including generative AI, partnering with school districts, universities, nonprofit organizations, media and tech developers, and others in iterative development and evaluation. As the PI of a NSF-funded grant, she is studying the use of generative AI in undergraduate writing courses. She also studies secondary student writing as a member of the IES-funded national WRITE Center. She received her B.A. in English and her Ph.D. in Education at U.C. Irvine and her J.D. at U.C. Berkeley.

Editorial Board

Stephen Aguilar, *University of Southern California, US*

Maha Bali, *American University in Cairo, Egypt*

Irene-Angelica Chounta, *University of Duisburg-Essen, Germany*

Shayan Doroudi, *University of California, Irvine, US*

María Florencia Ripani, *Ceibal Foundation, Uruguay*

Bart Rientes, *The Open University, UK*

Neil Selwyn, *Monash University, Australia*

Jiahong Su, *The University of Hong Kong*

Ulrich Trautwein, *University of Tübingen, Germany*

Ying Xu, *Harvard University, US*

About the Series

Generative AI is one of the most disruptive technologies in modern history, with the potential to dramatically transform education for better or worse. This series will address cutting-edge topics on the intersection of generative AI with educational research and practice for diverse learners from early childhood to adult.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914

Cambridge Elements

Generative AI in Education

Elements in the Series

Generative AI in Computer Science Education
Diana Franklin, Paul Denny, David A. Gonzalez-Maldonado and Minh Tran

Generative Artificial Intelligence and Language Teaching
Benjamin Luke Moorhouse and Kevin M. Wong

Academic Integrity in the Age of AI
Tricia Bertram Gallant, Mary Davis and Zeenath Reza Khan

Growing up with Artificial Intelligence: Implications for Child Development
Ying Xu and Kunlei He

A full series listing is available at: www.cambridge.org/EAIE

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 02:08:46, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/9781009503914