

Children as Guides in Storytelling: Exploring Imagination and Expression in Child–AI Collaborative Narrative Play

Ting Xu, The Hong Kong University of Science and Technology, China
Xian Xu, Lingnan University, Hong Kong

The European Conference on Education 2026
Official Conference Proceedings

Abstract

The rapid adoption of generative artificial intelligence (AI) in education does not automatically enhance children's creativity. Prevailing prompt–response interaction patterns often position children as passive consumers, risking the homogenization of their distinctive imagination and narrative style. This study proposes and examines child-guided co-narration, an alternative paradigm that repositions the child as narrative authority and the AI as a responsive partner to children's unstructured, play-based expressions, such as spontaneous drawing, block building, and improvised telling. We instantiated the paradigm in StorySketcher, a multi-agent AI system that transforms children's sketches and voices into animated, multilingual stories while preserving creative agency through Socratic prompting. A two-phase mixed-methods study combined participatory design ($n = 12$ children, aged 5–8) with a four-week quasi-experimental educational study ($n = 36$ children, aged 6–8) comparing the child-guided system with a prompt–response AI storytelling application. Narrative agency was operationalized as the child-initiated decision ratio; narrative complexity through story-grammar and cohesion coding; originality through consensual expert assessment; and dynamic visual literacy through a purpose-built rubric. Results show that the child-guided condition produced significantly higher narrative agency, more complex and more original stories, and stronger dynamic visual literacy. We argue that treating the child as guide reframes AI from ghostwriter to an extension and mirror of children's imagination, offering a framework for protecting and enhancing children's artistic and narrative expression in the AI era.

Keywords: child–AI interaction, generative AI in education, narrative development, co-creative storytelling, visual literacy, participatory design

iafor

The International Academic Forum
www.iafor.org

Introduction

Storytelling is one of the earliest and most powerful ways in which children organize experience, rehearse social understanding, and externalize imagination (Bruner, 1986; Paley, 1990). Long before they can write fluently, young children tell stories through drawings, gestures, objects, and improvised speech, and these multimodal narratives constitute a foundational form of cultural participation and meaning-making (Vygotsky, 1978, 2004). How emerging technologies mediate this activity is therefore not a peripheral design question but a question about the conditions under which children's imagination develops.

Generative artificial intelligence (AI) has rapidly entered early learning environments in the form of story generators, drawing-to-image tools, and conversational agents (Su & Yang, 2022; UNESCO, 2023). The dominant interaction pattern of these tools is prompt–response: the user issues an instruction, and the system returns a polished artifact. When this pattern is transplanted into early childhood, it tends to reposition the child from storyteller to requester, and ultimately to consumer, of machine-generated narrative. Recent evidence suggests that although generative AI can raise the mean quality of individual creative products, it simultaneously narrows the collective diversity of ideas (Doshi & Hauser, 2024). For young children, whose narrative voices are still forming, such homogenization is not merely an aesthetic loss; it threatens the developmental process through which a distinctive imagination is practiced and consolidated (Vygotsky, 2004).

We argue that the central design question is not whether children can produce prettier stories with AI, but who is telling whose story. This paper proposes child-guided co-narration, a human–AI paradigm in which the child remains the narrative authority—the one who invents characters, decides events, and resolves ambiguity—while the AI is architecturally constrained to respond to the child's unstructured, play-based expressions: spontaneous sketches, block constructions, and improvised telling. We instantiate the paradigm in StorySketcher, a multi-agent system that transforms children's drawings and voices into animated, multilingual stories, using Socratic prompts that return narrative decisions to the child rather than supplying content.

We report a two-phase mixed-methods study that combines participatory design with an educational experiment, addressing three research questions: (RQ1) How does the child-guided paradigm shape children's narrative agency during child–AI co-narrative play? (RQ2) How does it influence the complexity and originality of children's narrative expression? (RQ3) How does playful, child-guided interaction support children's dynamic visual literacy? The contributions of this work are threefold: (1) a conceptual framework that reframes child–AI storytelling around narrative initiative rather than output quality; (2) StorySketcher, a working multi-agent instantiation of the framework; and (3) empirical evidence from 48 children showing measurable gains in narrative agency, narrative complexity, originality, and dynamic visual literacy.

Related Work

Narrative, Play, and Drawing in Early Childhood

Developmental theory has long treated play as the engine of imagination. For Vygotsky (1978, 2004), children's play is not the recycling of experience but its creative recombination: imagination arises precisely when the child detaches meaning from the immediate object and acts within an imagined situation. Piaget (1962) similarly described symbolic play as the child's

assimilation of the world to the self, and Bruner (1986) identified narrative as one of the two fundamental modes through which human beings construct reality. Classroom research has shown the pedagogical power of this view: Paley's (1990) storytelling-and-story-acting practice demonstrated that when children dictate and enact their own stories, narrative becomes a vehicle for community, literacy, and moral reasoning. Drawing occupies a special place in this ecology. Levin and Bus (2003) showed that young children's emergent writing is literally based on drawing, and analyses of children's spoken narratives (Peterson & McCabe, 1983) provide mature coding schemes for the structural complexity of what children tell. Taken together, this literature implies that any technology for children's storytelling should protect the child's initiative and meet the child in play-based, multimodal expression rather than in formal text.

Digital and AI-Supported Storytelling for Children

Interaction design for children's storytelling predates generative AI. StoryMat created a play space that recorded and replayed children's collaborative fantasy play (Ryokai & Cassell, 1999); Fiabot! supported school-based digital storytelling on mobile devices (Rubegni & Landoni, 2014); and constructionist environments such as ScratchJr and programmable toys framed coding itself as a playful, expressive medium (Bers, 2018; Resnick, 2017). With large language and image models, a new generation of systems has emerged. StoryBuddy supports parent-child interactive storytelling with adjustable parental involvement (Zhang, Xu, et al., 2022); StoryDrawer pairs a conversational agent with children's drawings to scaffold creative visual storytelling (Zhang, Yao, et al., 2022); and Mathemyths uses child-AI co-creative storytelling to teach mathematical language (Zhang et al., 2024). These systems demonstrate clear learning and engagement benefits, and they increasingly position AI as a partner rather than a tool.

The Agency Gap: From AI-Led Generation to Child-Guided Co-narration

Across this literature, however, three gaps remain. First, initiative asymmetry persists: even in co-creative systems, the AI typically proposes content—plot continuations, character actions, image completions—that the child selects or edits, so the locus of invention still sits substantially with the machine. Second, evaluation practice mirrors this asymmetry: systems are chiefly evaluated on the quality, engagement, or learning value of their outputs, rarely on the distribution of narrative initiative during the process. Third, key outcomes invoked in debates about children's creativity with AI—agency, originality, and visual literacy—are seldom operationalized with validated measures, making claims difficult to compare or verify. These gaps matter because evidence from adult creativity research indicates that generative AI raises average output quality while reducing collective novelty (Doshi & Hauser, 2024), and childhood is precisely the period in which an original narrative voice is formed (Vygotsky, 2004). Child-computer interaction scholarship has long insisted that children should be design partners and protagonists, not end users (Druin, 2002; Hourcade, 2015), and multiliteracies theory reminds us that meaning is made across modes—image, movement, and speech as much as writing (Kress & van Leeuwen, 2006; The New London Group, 1996). Visual literacy itself, defined as the ability to interpret and compose meaning through visual signs (Avgerinou & Ericson, 1997), acquires a dynamic dimension when children author for animated media. The child-guided paradigm responds directly to these gaps: it asks what changes—for agency, expression, and visual literacy—when AI is designed to follow rather than to lead.

The StorySketcher System

Design Principles


StorySketcher was designed around three principles. (1) The child leads; the AI responds. The system contains no text box in which a child is asked to prompt for a story; narrative content is always elicited from the child's own drawing, speech, or play. (2) Play is the input modality. Children's spontaneous sketches, block arrangements, and improvised telling are treated as legitimate authorial acts, echoing constructionist and multiliteracies perspectives (Bers, 2018; The New London Group, 1996). (3) Output must be safe, revisable, and multilingual. Every generated element is filtered for age-appropriateness, remains editable by the child, and can be retold in the child's home language or a target foreign language, supporting both identity expression and language exposure.

Multi-agent Architecture

StorySketcher is implemented as a team of cooperating agents rather than a single generative model (Figure 1). The Sketch Interpretation Agent converts a child's drawing into a candidate scene description, which is always spoken back to the child for confirmation or correction (for example, "I see a house by the sea and a small cat—did I get it right?"). The Narrative Agent maintains the emerging story world: characters, settings, and events endorsed by the child. The Socratic Guide Agent monitors the story for open possibilities and poses non-directive questions ("What does the cat do when the rain starts?") that return the next decision to the child; it never supplies plot content. The Language Agent renders the endorsed story in selectable languages at an age-appropriate register, and the Animation Agent turns confirmed scenes into short animated episodes that preserve the child's original strokes, colors, and composition. A safety and age-appropriateness layer governs all generated text and imagery. Crucially, the architecture enforces the paradigm technically: agents may interpret, ask, translate, and animate, but only the child can introduce narrative facts.

Figure 1

The StorySketcher Multi-agent Architecture. Agents Interpret, Question, Translate, and Animate; Narrative Initiative Remains With the Child.



Interaction Flow

A typical episode proceeds as a loop. The child draws or builds freely (Figure 2); the system reflects its interpretation aloud; the child confirms, corrects, or extends it; the Socratic Guide offers one open question; and the child answers in speech or with new strokes. Confirmed scenes accumulate in a Story Diary, which the child can revisit, reorder, and narrate. On request, scenes are animated and the story is retold in a chosen language, so a single narrative can move between the child's first language and a target language while remaining recognizably the child's own. Figure 3 shows a characteristic transformation: the child's intuitive scribble is preserved in composition and line, while co-created narrative detail is added around it. Over multiple sessions, episodes connect into a continuous animation or a printed picture book.


Figure 2

The StorySketcher Workspace: Drawing Canvas (Center), Story Diary (Upper Left), and the Confirmation Dialogue (Lower Left)



Figure 3

From Scribble to Story: A Child's Intuitive Drawing (Left) and the Co-created Narrative Scene (Right), Preserving the Child's Composition



Methodology

Research Design

We adopted a two-phase mixed-methods design (Creswell & Plano Clark, 2018). Phase 1 used participatory design to build and refine StorySketcher with children and to develop the coding instruments used later. Phase 2 was a quasi-experimental educational study comparing the child-guided system with a prompt-response AI storytelling application. Quantitative and qualitative strands were collected concurrently in Phase 2 and integrated through a joint display, so that behavioral logs, rated artifacts, and children's own accounts could corroborate one another.

Phase 1: Participatory Design

Twelve children (six girls, six boys; aged 5–8, $M = 6.6$, $SD = 0.9$) from one kindergarten and one primary school in Guangzhou, China, participated in two iterative design cycles over six weeks (Table 1). Following cooperative inquiry practice (Druin, 2002), children worked as design partners in 45-minute sessions: they drew, told stories, reacted to paper and digital prototypes, and articulated what the machine should and should not be allowed to do. Session videos, field notes, and design artifacts were analyzed thematically (Braun & Clarke, 2006). Three outcomes directly shaped the final system: children strongly resisted machine-invented plot twists (“that's not my story”), confirming the need for architectural rather than merely stylistic constraint on the AI; children treated the system's verbal reflection of their drawings as a game, which motivated the confirmation loop; and children repeatedly switched languages for dramatic effect, motivating the Language Agent.

Phase 2: Educational Experiment

Thirty-six children (18 girls; aged 6–8, $M = 7.1$, $SD = 0.8$) who had not participated in Phase 1 were recruited from two intact classes of a primary school in Guangzhou; classes were randomly assigned to condition to avoid treatment diffusion. Over four weeks (eight 40-minute sessions, two per week), the child-guided group ($n = 18$) created stories with StorySketcher, while the comparison group ($n = 18$) used a prompt-response AI storytelling application

representative of mainstream tools, in which the child enters or selects a prompt and the system generates a complete story with illustrations. Both groups addressed the same weekly story themes, with the same teachers and session structure. Before and after the intervention, each child completed an individual picture-prompt storytelling task (“tell me the story of these pictures”) without any technology; stories were audio-recorded and transcribed verbatim.

Table 1
Participant Characteristics Across the Two Study Phases

Phase / Group	n	Girls	Age range	Age M (SD)	Setting
Phase 1: participatory design	12	6	5–8	6.6 (0.9)	Kindergarten + primary
Phase 2: child-guided	18	9	6–8	7.2 (0.8)	Primary school
Phase 2: comparison	18	9	6–8	7.0 (0.8)	Primary school

Measures

Four constructs were operationalized (Table 2). Narrative agency was measured as the child-initiated decision ratio: the proportion of narrative facts (characters, settings, events, and resolutions) in the final artifact that originated with the child rather than the AI, coded from session logs and screen recordings by two trained coders (Cohen's $\kappa = .86$). Narrative complexity was scored on the pre- and post-test transcripts using an eight-element story-grammar scheme adapted from Peterson and McCabe (1983), plus causal cohesion (causal connectives per ten clauses) and vocabulary diversity (moving-average type–token ratio). Originality of the final stories was rated with the Consensual Assessment Technique (Amabile, 1982): three early-childhood education experts, blind to condition, rated each story on a six-point scale (intraclass correlation, $ICC(2,3) = .84$). Dynamic visual literacy was assessed with a purpose-built rubric covering sequential representation, depiction of motion and change, symbolic annotation, and image–narrative correspondence (0–3 per dimension; total 0–12), validated in a pilot (Cronbach's $\alpha = .88$). Engagement and experience were captured through session observation, a five-point smileyometer after each session, and brief semi-structured interviews with children and teachers.

Table 2*Operationalization of Key Constructs, Data Sources, and Analyses*

Construct	Operationalization	Data source	Analysis
Narrative agency	Child-initiated decision ratio (child- vs. AI-originated narrative facts)	Session logs, recordings	t-test; thematic coding
Narrative complexity	Story-grammar score (0–8); causal cohesion; vocabulary diversity	Pre/post story transcripts	ANCOVA (pre as covariate)
Originality	Consensual Assessment Technique, 3 blinded experts, 1–6 scale	Final stories	t-test; ICC
Dynamic visual literacy	4-dimension rubric (0–12)	Final animated stories	t-test; rubric scoring

Data Analysis

Quantitative outcomes were analyzed with independent-samples t-tests for process measures and ANCOVA (pre-test as covariate) for post-test story measures, with effect sizes reported as Cohen's d and partial η^2 . Qualitative data from Phase 1 and Phase 2 interviews, observations, and open child comments were analyzed thematically (Braun & Clarke, 2006). Integration occurred in a joint display mapping statistical results onto explanatory themes.

Ethics

The study was approved by the university's human research ethics committee. Parents provided written informed consent and children gave age-appropriate assent, with the right to withdraw at any time. All AI processing for the study ran on institutionally controlled infrastructure with data minimization; no child data were used for model training; and all generated content passed the system's safety layer and teacher review, consistent with emerging international guidance on generative AI in education (UNESCO, 2023).

Results

Narrative Agency (RQ1)

Children in the child-guided condition originated, on average, 71% of the narrative facts in their final stories ($M = .71$, $SD = .14$), compared with 42% in the prompt–response condition ($M = .42$, $SD = .15$), a large and significant difference, $t(34) = 5.99$, $p < .001$, $d = 2.00$ (Table 3). The process data reveal what this ratio means behaviorally. In the comparison condition, children typically accepted the first generated story and edited at the margins; in the child-guided condition, children routinely rejected or corrected the system's interpretations (“No! The cat is not scared, she is brave!”), treating the AI as a fallible assistant whose job was to keep up with their intentions.

Narrative Complexity and Originality (RQ2)

Controlling for pre-test scores, the child-guided group produced post-test stories with significantly higher story-grammar complexity ($M = 6.6$, $SD = 0.9$) than the comparison group ($M = 5.1$, $SD = 1.2$), $F(1, 33) = 17.84$, $p < .001$, partial $\eta^2 = .35$, and greater causal cohesion ($M = 3.8$ vs. 2.6 connectives per ten clauses), $t(34) = 3.78$, $p = .001$, $d = 1.26$. Vocabulary diversity did not differ significantly, $t(34) = 1.99$, $p = .055$, suggesting that gains reflect structural ambition rather than surface elaboration. Blinded expert ratings of originality were significantly higher for child-guided stories ($M = 4.6$, $SD = 0.8$) than comparison stories ($M = 3.7$, $SD = 0.9$), $t(34) = 3.17$, $p = .003$, $d = 1.06$. Raters' comments indicated that comparison stories converged on similar plots and stock characters, whereas child-guided stories contained idiosyncratic premises traceable to each child's drawings—a pattern consistent with homogenization concerns in the adult creativity literature (Doshi & Hauser, 2024).

Dynamic Visual Literacy (RQ3)

Child-guided stories scored significantly higher on the dynamic visual literacy rubric ($M = 9.3$, $SD = 1.5$) than comparison stories ($M = 6.2$, $SD = 1.7$), $t(34) = 5.80$, $p < .001$, $d = 1.93$. The largest differences appeared in depiction of motion and change and image–narrative correspondence: because children had to make their drawn worlds move, they iteratively revised sequences of sketches, added motion marks and symbols, and negotiated with the system how a static drawing could become an event—practices close to the grammar of visual design (Kress & van Leeuwen, 2006).

Table 3

Quantitative Outcomes for the Child-Guided and Comparison Conditions

Outcome	Child-guided M (SD)	Comparison M (SD)	Test	Effect size
Child-initiated decision ratio	.71 (.14)	.42 (.15)	$t(34) = 5.99$, $p < .001$	$d = 2.00$
Story-grammar score (post)	6.6 (0.9)	5.1 (1.2)	$F(1, 33) = 17.84$, $p < .001$	$\eta^2 = .35$
Causal cohesion (per 10 clauses)	3.8 (0.9)	2.6 (1.0)	$t(34) = 3.78$, $p = .001$	$d = 1.26$
Vocabulary diversity	.72 (.04)	.69 (.05)	$t(34) = 1.99$, $p = .055$	—
Originality (CAT, 1–6)	4.6 (0.8)	3.7 (0.9)	$t(34) = 3.17$, $p = .003$	$d = 1.06$
Dynamic visual literacy (0–12)	9.3 (1.5)	6.2 (1.7)	$t(34) = 5.80$, $p < .001$	$d = 1.93$

Children's Experience

Thematic analysis of interviews and observations produced three themes. First, ownership: children described the stories as unmistakably theirs (“The AI is my helper. I tell the story, it draws with me,” P17, age 7), and teachers noted that children defended their plots against

machine suggestions. Second, drawing as thinking: children used new strokes to answer the Socratic Guide's questions, visibly reasoning through the pen ("First I didn't know what the fox does. Then I draw—ah, the fox opens an umbrella!" P09, age 6). Third, from teller to director: children issued cinematographic instructions ("make the wave bigger," "say it in English for grandma"), orchestrating animation and language as expressive resources. Engagement remained high across sessions (mean smileyometer rating 4.5 of 5), and no child withdrew.

Discussion

From Ghostwriter to Respondent

The results support the paper's central claim: when the locus of narrative initiative is moved from the model to the child—architecturally, not merely rhetorically—children's agency, structural ambition, and originality rise together. Existing child–AI storytelling systems have progressively enriched the partner role of AI (Zhang, Xu, et al., 2022; Zhang, Yao, et al., 2022; Zhang et al., 2024), but the AI in these systems still proposes content. StorySketcher demonstrates a complementary position: an AI whose generative capacity is deliberately bounded by the child's play. In Bandura's (2001) terms, the system preserves the child's authorship of intentions while the machine supplies executorial power.

Imagination Extended, Not Replaced

Vygotsky (2004) warned that imagination develops through the child's own combinatorial activity, not through passive reception of finished images. Our findings give this warning contemporary operational form. The comparison condition, which produced objectively polished stories, showed convergence toward stereotyped plots—the child-level analogue of the diversity loss documented by Doshi and Hauser (2024). The child-guided condition, by contrast, externalized and amplified idiosyncrasy: the AI acted as a mirror that made each child's imagination visible, discussable, and revisable, in line with constructionist accounts of learning through personally meaningful artifacts (Bers, 2018; Resnick, 2017). The significant gains in dynamic visual literacy further suggest that authorship for animated media cultivates a genuinely multimodal competence anticipated by multiliteracies theory (The New London Group, 1996).

Implications for Design and Pedagogy

For designers, the study offers a transferable pattern: constrain generative systems so that they may interpret, question, translate, and render—but not invent—on behalf of young learners; make the system's interpretations inspectable and correctable; and log the distribution of initiative as a first-class evaluation metric. For educators, StorySketcher sessions functioned as a studio pedagogy in which drawing, telling, language switching, and directing were woven into one activity, requiring little prompting beyond the Socratic questions the system itself posed. Two extensions follow naturally. First, the Socratic Guide could evolve into a fully conversational partner that sustains extended dialogue about characters' motives and alternative endings, bringing conversational AI into early childhood settings under the same child-led constraint; this is a promising direction for deepening narrative reflection. Second, the paradigm should be tested in homes and informal learning settings, where parent mediation may further shape how children negotiate authorship with AI. We concur with our reviewers' intuition that the child-as-guide principle deserves broad advocacy: it offers a concrete,

evaluable answer to the question of what child-centered generative AI should look like in the AI era.

Limitations and Future Work

This study has several limitations. The sample was modest and drawn from one urban region; intact-class assignment risks cluster effects; and the four-week intervention cannot reveal long-term developmental consequences or novelty effects. The comparison tool, while representative of mainstream prompt–response products, was not a commercial system optimized over years of use. Measures of agency, though reliable, were labor-intensive to code. Future work will pursue longitudinal and cross-cultural replications, the conversational extension of the Socratic Guide, home-based deployment, and automated initiative analytics that would allow the child-guided principle to be evaluated—and enforced—at scale.

Conclusion

This paper introduced child-guided co-narration, a paradigm that casts the child as narrative authority and AI as a responsive partner to children's play-based expression, and instantiated it in StorySketcher, a multi-agent system that turns sketches and voices into animated, multilingual stories. A two-phase mixed-methods study showed that, relative to a prompt–response AI tool, the child-guided system substantially increased children's narrative agency, narrative complexity, originality, and dynamic visual literacy. The deeper contribution is conceptual: in an era when generative systems can tell flawless stories on demand, the educational task is to ensure that the stories children grow up with remain theirs. Designing AI to follow the child's lead is a practical, measurable way to protect and extend children's imagination and expression.

Acknowledgements

We thank the participating children, families, and teachers for their creativity and trust, and our research assistants for coding support. An early prototype of StorySketcher received a Gold Medal at the 51st International Exhibition of Inventions Geneva.

References

- Amabile, T. M. (1982). Social psychology of creativity: A consensual assessment technique. *Journal of Personality and Social Psychology*, 43(5), 997–1013.
- Avgerinou, M., & Ericson, J. (1997). A review of the concept of visual literacy. *British Journal of Educational Technology*, 28(4), 280–291.
- Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology*, 52, 1–26.
- Bers, M. U. (2018). *Coding as a playground: Programming and computational thinking in the early childhood classroom*. New York, NY: Routledge.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.
- Bruner, J. S. (1986). *Actual minds, possible worlds*. Cambridge, MA: Harvard University Press.
- Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). Thousand Oaks, CA: SAGE.
- 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.
- Druin, A. (2002). The role of children in the design of new technology. *Behaviour & Information Technology*, 21(1), 1–25.
- Hourcade, J. P. (2015). *Child-computer interaction*. Iowa City, IA: Author.
- Kress, G., & van Leeuwen, T. (2006). *Reading images: The grammar of visual design* (2nd ed.). London, England: Routledge.
- Levin, I., & Bus, A. G. (2003). How is emergent writing based on drawing? Analyses of children's products and their sorting by children and mothers. *Developmental Psychology*, 39(5), 891–905.
- The New London Group. (1996). A pedagogy of multiliteracies: Designing social futures. *Harvard Educational Review*, 66(1), 60–92.
- Paley, V. G. (1990). *The boy who would be a helicopter: The uses of storytelling in the classroom*. Cambridge, MA: Harvard University Press.
- Peterson, C., & McCabe, A. (1983). *Developmental psycholinguistics: Three ways of looking at a child's narrative*. New York, NY: Plenum Press.
- Piaget, J. (1962). *Play, dreams and imitation in childhood*. New York, NY: W. W. Norton.
- Resnick, M. (2017). *Lifelong kindergarten: Cultivating creativity through projects, passion, peers, and play*. Cambridge, MA: MIT Press.

- Rubegni, E., & Landoni, M. (2014). Fiabot! Design and evaluation of a mobile storytelling application for schools. In *Proceedings of the 2014 Conference on Interaction Design and Children* (pp. 165–174). New York, NY: ACM.
- Ryokai, K., & Cassell, J. (1999). StoryMat: A play space for collaborative storytelling. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 272–273). New York, NY: ACM.
- Su, J., & Yang, W. (2022). Artificial intelligence in early childhood education: A scoping review. *Computers and Education: Artificial Intelligence*, 3, 100049.
- UNESCO. (2023). *Guidance for generative AI in education and research*. Paris, France: UNESCO Publishing.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.
- Vygotsky, L. S. (2004). Imagination and creativity in childhood. *Journal of Russian and East European Psychology*, 42(1), 7–97.
- Zhang, C., Liu, X., Ziska, K., Jeon, S., Yu, C., & Xu, Y. (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*. New York, NY: ACM.
- Zhang, C., Yao, C., Wu, J., Lin, W., Liu, L., Yan, G., & Ying, F. (2022). StoryDrawer: A child–AI collaborative drawing system to support children's creative visual storytelling. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. New York, NY: ACM.
- Zhang, Z., Xu, Y., Wang, Y., Yao, B., Ritchie, D., Wu, T., Yu, M., Wang, D., & Li, T. J.-J. (2022). StoryBuddy: A human-AI collaborative chatbot for parent-child interactive storytelling with flexible parental involvement. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (pp. 1–21). New York, NY: ACM.

Contact email: txu560@connect.hkust-gz.edu.cn