BUILDING BLOCKS FOR UNDERSTANDING ARTIFICIAL INTELLIGENCE: DESIGNING INTERACTIVE AI LEARNING EXPERIENCES FOR YOUNG CHILDREN AND THEIR FAMILIES IN A MUSEUM SETTING

E. Greenwald[1], D. Cavero[1], K. Grindstaff[1], A. Krakowski[1], T. Hurt[1], N. Wang[2]
[1]*University of California, Berkeley (UNITED STATES)*
[2]*University of Southern California (UNITED STATES)*

Abstract

As artificial intelligence (AI) profoundly reshapes our personal and professional lives, there are growing calls to support pre-college aged youth as they develop capacity to engage critically and productively with AI. While efforts to introduce AI concepts to pre-college aged youth have largely focused on older teens, there is growing recognition of the importance of developing AI literacy among younger children. Today's youth already encounter and use AI regularly, but they might not yet be aware of its role, limitations, risks, or purpose in a particular encounter, and may not be positioned to question whether it should be doing what it's doing. In response to this critical moment to develop AI learning experiences that can support children at this age, researchers and learning designers at the University of California's Lawrence Hall of Science, in collaboration with AI developers at the University of Southern California's Institute for Creative Technologies, have been iteratively developing and studying a series of interactive learning experiences for public science centers and similar out-of-school settings. The project is funded through a grant by the National Science Foundation and the resulting exhibit, The Virtually Human Experience (VHX), represents one of the first interactive museum exhibits in the United States designed explicitly to support young children and their families in developing understanding of AI. The coordinated experiences in VHX include both digital (computer-based) and non-digital ("unplugged") activities designed to engage children (ages 7-12) and their families in learning about AI. In this paper, we describe emerging insights from a series of case studies that track small groups of museum visitors (e.g. a parent and two children) as they interact with the exhibit. The case studies reveal opportunities and challenges associated with designing AI learning experiences for young children in a free-choice environment like a public science center. In particular, we focus on three themes emerging from our analyses of case data: 1) relationships between design elements and collaborative discourse within intergenerational groups (i.e., families and other adult-child pairings); 2) relationships between design elements and impromptu visitor experimentation within the exhibit space; and 3) challenges in designing activities with a low threshold for initial engagement such that even the youngest visitors can engage meaningfully with the activity. Findings from this study are directly relevant to support researchers and learning designers engaged in rapidly expanding efforts to develop AI learning opportunities for youth, and are likely to be of interest to a broad range of researchers, designers, and practitioners as society encounters this transformative technology and its applications become increasingly integral to how we live and work.

Keywords: Artificial Intelligence, education, museums.

1 INTRODUCTION

The increasing presence and importance of Artificial Intelligence (AI) in our society has led to calls for its inclusion at all levels of education [1]. However, the field is only beginning to understand what makes AI learning experiences most effective and developmentally appropriate, especially for young children. In response to this need, researchers and learning designers are developing an increasing range of AI learning experiences and curricula for young children. Much of this work has been based on AI learning goals that were developed to be developmentally appropriate for each grade band within K–12 education [2]. However, research on how best to support younger children's interactions with technology to promote AI literacy is still nascent and primarily focused on facilitated learning environments, whether formal, like a classroom, or informal, like an afterschool program. Given the important role of public science museums in the STEM learning ecosystem, especially for young children [3, 4], we set out to investigate how young learners interact with and make sense of AI systems in free choice environments, and what strategies best position them to come away with new understanding of these systems.

For this effort, researchers and learning designers at the University of California's Lawrence Hall of Science, in collaboration with AI developers at the University of Southern California's Institute for Creative Technologies, have been iteratively developing and studying a series of interactive learning experiences for public science centers and similar out-of-school settings. The project is funded through a grant by the National Science Foundation and the resulting exhibit, The Virtually Human Experience (VHX), represents one of the first interactive museum exhibits in the United States designed explicitly to support young children and their families in developing understanding of AI. The coordinated experiences in VHX include both digital (computer-based) and non-digital ("unplugged") activities designed to engage children (ages 7-12) and their families in learning about AI. In this paper, we describe emerging insights from a series of case studies that track small groups of museum visitors (e.g. a parent and two children) as they interact with the exhibit. The case studies reveal opportunities and challenges associated with designing AI learning experiences for young children in a free-choice environment like a public science center. In particular, we focus on three themes emerging from our analyses of case data: 1) relationships between design elements and collaborative discourse within intergenerational groups (i.e., families and other adult-child pairings); 2) relationships between design elements and impromptu visitor experimentation within the exhibit space; and 3) challenges in designing activities with a low threshold for initial engagement such that even the youngest visitors can engage meaningfully with the activity.

2 RELATED WORK

AI learning among young children. While there is a growing body of work contributing to interaction design and research of AI learning experiences for pre-college aged youth, there are fewer studies that explore how younger children (ages 4-10) make sense of AI [5]. Recent studies have demonstrated potential for young children's interaction with social robots [6] and conversational agents [7] to support their development of AI literacy and understanding of AI models through a set of classroom-based instructional activities. Evidence from recent work [8] points to engagement with human emotional expression as a promising approach for supporting young children in AI learning, because making sense of cues about human emotion is a ubiquitous and central element of their ongoing, developing socialization and emotional regulation [9, 10]. This approach dovetails with recommendations in the AI4K12 "Big Ideas" framework [11] that young children in grades K-2 (approximately ages 5-8) should investigate AI models that recognize expressions and inferred emotional states in human faces to make sense of how AI agents interact with humans (referred to as "Big Idea 4: Natural interaction" in the AI4K12 framework).

Learning AI in Informal Spaces. Designed spaces for informal science education such as science centers and museums can promote science learning and positive affective engagement with science in free-choice environments [12, 13, 14,15], complementing and extending more constrained school-based opportunities for science learning [16]. Informal learning contexts can therefore serve as rich environments to conduct research about how children build understanding about concepts outside of standards to which formal learning contexts are generally bound, including concepts related to AI literacies. For example, Duri Long and colleagues have articulated design principles for AI literacy [1] and investigated AI learning experiences designed for museum settings, providing evidence for the strategies of embodiment and opportunities for metacognition to support the development of AI literacy [17, 18]. Designing experiences in science centers and museums necessitates particular attention to dispositional components of learning, such as engagement, curiosity, and fascination [19]. Such research can therefore contribute to the field's broader understanding of how interaction design approaches can promote children's interest and engagement with AI and provide traction for conceptual learning that translates across learning contexts. Yet, there is a dearth of research being conducted about how to promote AI literacy and engagement in informal learning spaces [20].

3 EXHIBIT DESIGN AND PEDAGOGICAL APPROACH

The VHX project has enabled us to explore how youth make sense of AI and, in particular, the interplay between a child's developing mental model of how an AI system operates, and the child's own metacognitive understanding of how they perform similar tasks. Drawing on situative learning theories [21] the exhibit's design emerged through an iterative process of piloting and testing a range of strategies and activity structures to support youth sense-making about AI. Through this iterative process, we developed the following core pedagogical strategies that have shown evidence of promise for promoting youth engagement and understanding of AI and that served to ground our current exhibit design:

1. **Situate engagement with AI in the human experience.** The AI experiences in the Virtually Human exhibit are designed to explicitly connect to human experiences, so that children and their families interact with AI systems through contexts that are familiar and relatable. This approach positions children to investigate the capabilities of each AI model in the exhibit by drawing upon their own human experience and shared lived experience with their family members or peers interacting with the exhibit.
2. **Metacognitive Embodiment.** Through the iterative design and prototyping of non-digital, “unplugged” interactions that complement children’s interactions with AI virtual humans, we developed the strategy of metacognitive embodiment [22]. Metacognitive embodiment involves designing explicit opportunities for children to reflect upon their own process in performing a cognitive task, and leverage that to advance understanding of how an AI might perform a similar task.
3. **Virtual Human as facilitator.** To support engagement with the AI experiences and motivate curiosity about how they work, we position the Virtual Human character as a friendly facilitator. The Virtual Human introduces herself as an AI, and provides descriptive information about what the AI is doing in real time.

Figure 1. The Virtually Human Experience.

Aligned with these principles, the exhibit was organized around a series of paired digital and non-digital experiences for youth that can leverage the interplay between how youth reason through a task, such as recognizing whether someone is happy, and how an AI might accomplish a similar task. Each activity pair was framed around two guiding questions: How do YOU do X? How does an AI do X? As seen in Figure 2, the specific question varied based on the activity. While the full exhibit includes seven paired activities, the three sets presented below give a sense of how our pedagogical approach was implemented.

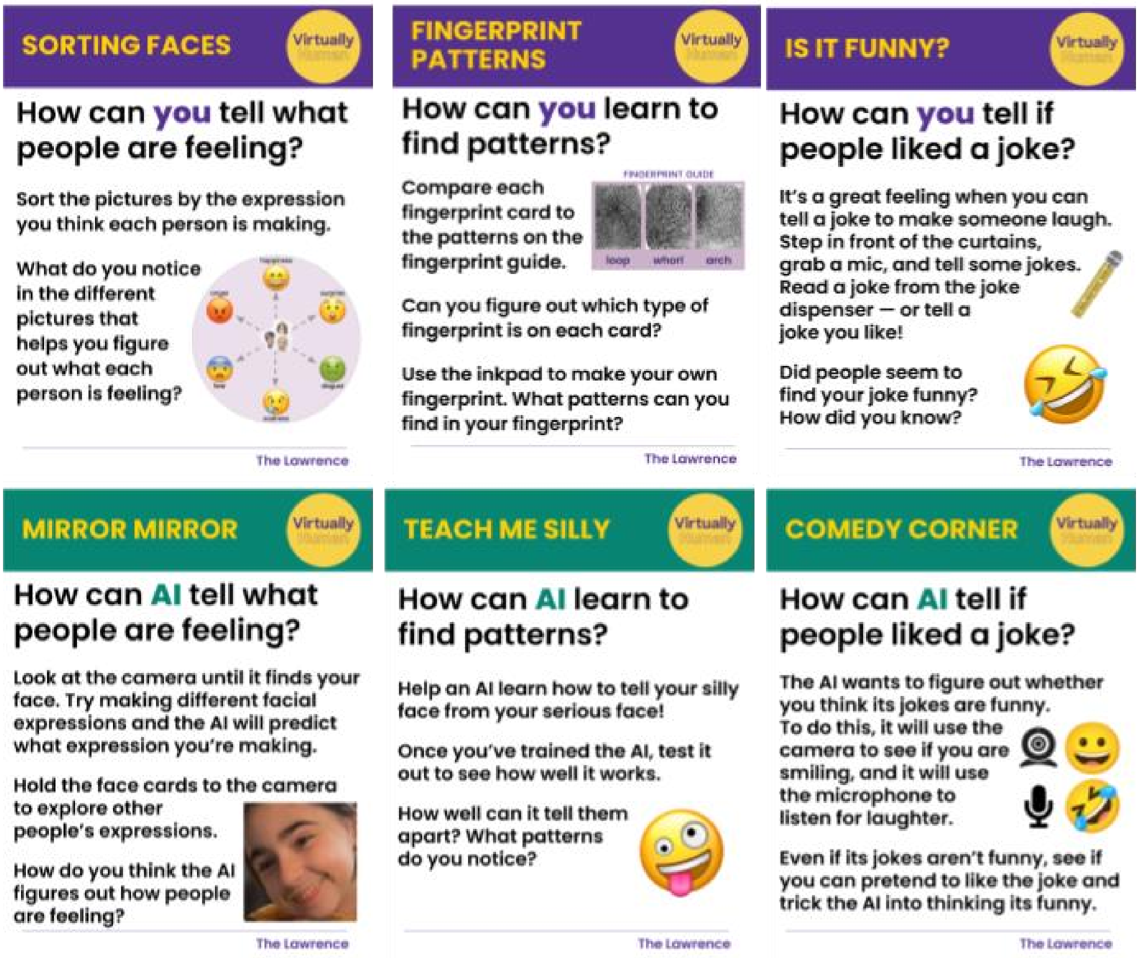

Figure 2. Interpretive Signs: Activity Descriptions.

4 METHODOLOGY

This manuscript focuses on a series of case studies of intergenerational groups of visitors (e.g., a child and parent, children and grandparents, or children and adult caretakers) as they enter into the exhibit space and interact with the various activities within it. Data for this study were collected over a three-day period at a public science center on the West Coast of the United States. The goal of data collection was to construct a series of case studies capturing the range of visitor experiences as they interact with an interactive exhibit about AI designed for children ages 6 and up. To construct each case study, researchers tracked the group of visitors and captured field observations, describing what activities the group engaged with, for how long, and what transpired at each activity station. The study was approved by our University's Institutional Review Board and visitors participating in the study completed consent forms upon entering the exhibit space. With a particular focus on intergenerational learning opportunities in the exhibit, fieldnotes captured discourse, behaviors, affect, and engagement within the group.

To support observational field notes, we conducted follow up interviews with visitors (adults and children), with the discussion mediated through a retrospective walk-through of the exhibit. For these retrospective walk-throughs, visitors were asked to return to particular activities (specifically, those featured in Figure 2) and describe their experience. This strategy enabled us to conduct the observational study in as naturalistic conditions as possible (i.e., without controlling the visitor's initial experience), while still capturing additional information about the visitor's internal experiences as they moved through the exhibit space. These follow up interviews were audio recorded and transcribed and synthesized with the observational notes to construct a full case study of for each visitor group. This work resulted in seven full case studies of groups with 3-7 visitors per group.

Analysis was grounded in the principles of multiple-case studies analysis [23, 24], in which we examined similarities and differences between cases and themes that cut across multiple cases. This also served as a data reduction step, such that we were able to identify a subset of cases that could serve to illustrate the most salient features observed across the full set. To clarify, the three cases described below are not composite cases, but cases selected from the larger set because they are particularly illustrative of what the research team observed throughout the study.

5 CASE SUMMARIES

As context for the case summaries, the exhibit room's interior for the interactive exhibit was designed specifically for the exhibit, and built on our design team's experience with visual design for young children: bright, child-friendly seating, large wall posters featured colorful, engaging graphics, with minimal text, and familiar symbols like emojis, and photos of children making the various expressions recognized by the Mirror Mirror activity. For example, the expression predictions at Mirror Mirror were conveyed using emojis, in addition to text labels. Further, at each activity station, was a tabletop interpretive sign with a description of what to do (Figure XX, above), and an explainer of the underlying AI (Figure 3, below). These interpretive signs were rarely a child's first point of engagement with the station's activity (they tended to jump right in and press whatever buttons were available to press), but they were often leveraged by adult caregivers, who would use the guiding questions to facilitate their child's experience.

Figure 3. Interpretive Signs: Activity Explainers.

While fully described case studies are outside the scope of this manuscript, we've summarized three cases, focusing on visitor experiences across two activities for each case. Together, the cases serve to illustrate themes discussed in the next section.

5.1 Case Study 1

Two children (around 6-9 years old) and their grandparents entered the Virtual Human room just after lunch. The young boy and his grandfather moved straight to Mirror Mirror (see Figure 2 for a summary of the activity structure), which was the closest activity to the door. As they started to explore the activity, the boy seemed a little bit hesitant and confused, so the grandfather gave instructions about how to do it, telling him "the camera's right over there" and "try the explore/tell me more mode." As they started to engage, the young girl and her grandmother came over and joined, reacting to the boy making facial expressions at the Virtual Human. At the station nearby were several laminated photographs of children with the various expressions the AI model could identify. The images made up a portion of the data set used to train the AI model. The siblings began to bring some of the photos over to the Mirror Mirror station, and took turns holding the images up to the camera and seemed delighted when the Virtual Human could recognize them.

Next, the boy transitioned into the Teach Me Silly activity. He laughed at his selection of silly and serious faces as he moved through the progression. When he got to the end the first time, the AI didn't do a great job of learning his silly and serious faces and classifying them. After a pause, it appeared that he had a moment of understanding about the exhibit and started to make serious faces both with and without glasses, testing if the AI could distinguish his face with and without glasses instead of silly and serious. It seemed to work well and he then tried it by turning his face left and turning his face right to train on face angle. This didn't work particularly well for him, not nearly as well as the glasses no glasses. The girl joined in and tried Teach Me Silly next. She asked the facilitator nearby to confirm it was silly and serious faces or another expression. It didn't work well the first time, so she tried again. During the portion of the experience where they ask if they are silly or serious faces, she pointed at which pictures matched the examples given, drawing visual lines in the air between them. She also did one round with her face looking left (during the "silly") and right (during the "serious" call), which worked much better than when the boy tried it. After spending time with 2 other activity stations, the group exited, having been in the room for about 1 hour.

5.2 Case Study 2

In the early afternoon (around 2pm), a dad entered with his 8-year-old twins. They are members of the museum and told us that they had come to the exhibit once before and that the girls wanted to come back to see if they could "trick" the Teach Me Silly station this time. As identical twins, they had a clear goal of trying to "trick" the AI. They first tried to train it with one of their faces as silly and one as serious to see if the Virtual Human could tell them apart. After it wasn't able to tell them apart, they tried making silly and serious faces but switching in halfway through to see what the AI would do. It seemed like the AI was able to determine silly from serious faces when the two twins switched in and out during the "live testing" phase of the Virtual Human experience. At the end of the experience, they focused on specific facial features that AI may have used to tell them apart. The dad told the girls, "You're trying to program AI to interact in a certain way." The group then tried another station about AI planning. However, they interacted with this only briefly (less than 2min) and both took a turn but were clearly not as engaged as with Teach Me Silly.

5.3 Case Study 3

Just before 11:00 on a Sunday morning, a mom and her 5-year-old daughter enter the Virtually Human exhibit. As most visitors do, based on its position relative to the entrance, the pair head to Mirror Mirror and spend a little under five minutes, with the girl making facial expressions for the VH to interpret. She likes seeing herself on the screen. Her mom provides space for the child to engage independently at this activity as the child appears to understand that the VH is interpreting her facial expression and showing with the display what emotion it "thinks" she is feeling. Afterward her mother commented that she felt her daughter "was able to kind of connect the dots of how AI is reading different pieces of her facial expressions and what that is translating to when she does it herself, which is exciting for her."

At each of the other activities in the exhibit, her mother takes on the role of facilitator, prompting her daughter with questions that elicit many of the noticings each activity is designed to elicit. For example, at Teach Me Silly, the VH is uncertain whether the girl's expression is silly or serious and asks, "Why do you think I struggled?" Her mother prompts the girl, asking "Do you think it's because your eyes weren't silly? And because you're mouth was silly? Were you half silly and half serious? Maybe that's why the VH didn't know which face you were making." Later when asked about her experience with that activity, and that moment that the VH had a hard time, the young girl said, "I think it's because like my eyes didn't show silliness, so [the VH] put it in a serious [category], but my mouth was kind of silly. My mouth was really, really, silly." She would not have come up with this explanation without her mother's prompting and facilitation, but the girl does understand that the VH is looking at different parts of her face to decide what expression she is making, even at the young age of five.

6 DISCUSSION: THEMES AND DESIGN CONSIDERATIONS

We discuss three themes (below) that have emerged thus far in our analysis of case data. We describe each and consider design features that may be contributing to them.

Reflection & projection, in high definition. It cannot be overstated, the power of a room filled with giant TVs upon which young children can project various images of themselves. Visitors of all ages enjoyed seeing themselves and their group on the big screens, playfully engaging with the stage and screen. This unquestionably contributed to overall engagement and dwell time, which is consistent with what we have observed since the earliest activity piloting. From a design perspective, we located the activities that projected images of the visitor toward the front of the room. Further, we elected to use a facial detection algorithm to “lock on” to the face of whoever is doing the activity, rather than just allowing the camera feed to capture whatever happens to be in it when the activity starts running. This decision helps in two ways. First, it makes it more likely that the data the AI uses is face data - which makes the prediction and classification systems more accurate. Second, it was clearly more fun for visitors to have their faces locked onto - the thrill of immediately seeing themselves seemed to promote sustained engagement.

Interactions were typically collaborative. We were surprised at how regularly parents were able to leverage signage and use prompting questions to facilitate their children’s sense-making about many of the activities. As the interpretive signs demonstrate (Figures 2 & 3), suggested prompts were brief and, across all activities, structured around a comparison between how the visitor might accomplish the task as a human, and how the AI might be accomplishing the same task. Parents typically facilitated by helping the child get started (e.g., attending to the camera, as in Case 1) and/or using questions similar to those in the interpretive signs (Case 3). Other design features that seemed to prompt collaborative include: *Redundancy*: for example, the VH says similar things in different ways to help learners pick up on what is being asked; *Show in addition to tell*: for example, we created images of the VH that give examples of silly and serious faces to help guide the activity flow and provide a model for the child to work from; *Ample pauses and wait time*: we chose a deliberate talking speed and added pauses to give learners time to comprehend what was being said.

Visitor experimentation can drive learning. One of the most compelling observations across cases was the range of impromptu experimentation visitors engaged in with the activities. Case 2 offers the example of twins trying to “trick” the AI through various experiments with the Teach Me Silly activity. Case 1 includes the example of visitors using material from one activity to test the AI system of another activity. From a design perspective, while the VH avatar asserts that one group of photos is Silly and the other is Serious, many learners realize that it is simply two groups and can therefore be whatever they want them to be (hat / no hat; me / my sister; etc). This experimentation, while not directly prompted for in the experience, is well-aligned with our goals and hopes for this experience. We also discovered an aspect of the setup for Teach Me Silly that seemed to elicit both confusion and also prompt further experimentation: Given that people enjoy seeing their photos be captured on the screen, they generally look to the side of the screen where their photos are appearing. For the “Silly” condition, this is on the left side of the screen. For “Serious”, it is on the right. Collectively, this means that learners often have a different head positioning for the Silly and Serious classes. This sometimes means that the model learns their head positioning in addition to or instead of the classes we intended for the model to learn. Sometimes this led to frustration (e.g., if the positioning pattern superseded the expression pattern, without them realizing it), but it sometimes prompted visitors to try again, and/or observe the images more closely to notice the actual pattern that the AI was picking up on to make its class determinations. Museum facilitators, and often parents, were instrumental in helping the transition from frustration and confusion to experimentation and discovery.

ACKNOWLEDGEMENTS

This research was supported by a grant from the National Science Foundation (DRL 2116109). Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

References [Arial, 10-point, left alignment, upper and lower case] should be cited according to the Bibliography and Citation Style https://iated.org/citation_guide

[1] Long, D., & Magerko, B. (2020, April). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1-16).

[2] Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019, July). Envisioning AI for K-12: What should every child know about AI?. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, No. 01, pp. 9795-9799).

[3] Allen, S., Campbell, P. B., Dierking, L. D., Flagg, B. N., Friedman, A. J., Garibay, C., & Ucko, D. A. (2008, February). Framework for evaluating impacts of informal science education projects.

In *Report from a National Science Foundation Workshop. The National Science Foundation, Division of Research on Learning in Formal and Informal Settings* (Vol. 73).

[4] Sullenger, K. (2006). Beyond School Walls: Informal Education and the Culture of Science. *Education Canada*, *46*(3), 15-18.

[5] 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.

[6] Williams, R., Park, H. W., & Breazeal, C. (2019, May). A is for artificial intelligence: the impact of artificial intelligence activities on young children's perceptions of robots. In *Proceedings of the 2019 CHI conference on human factors in computing systems* (pp. 1-11).

[7] Lin, P.; Van Brummelen, J.; Lukin, G.; Williams, R.; and Breazeal, C. 2020. Zhorai: Designing a conversational agent for children to explore machine learning concepts. In *Pro ceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 13381–13388.

[8] Vartiainen, H.; Tedre, M.; and Valtonen, T. 2020. Learning machine learning with very young children: Who is teaching whom? *International journal of child-computer interaction*, 25: 100182.

[9] Denham, S. A. 2006. Social-emotional competence as sup port for school readiness: What is it and how do we assess it? *Early education and development*, 17(1): 57–89.

[10] Denham, S. A. 2007. Dealing with feelings: how children negotiate the worlds of emotions and social relationships. *Cognitie, Creier, Comportament/Cognition, Brain, Behavior*, 11(1).

[11] Grover, S. (2024, March). Teaching AI to K-12 Learners: Lessons, Issues, and Guidance. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1* (pp. 422-428).

[12] Falk, J.; and Dierking, L. 2000. Learning from museums.

[13] Falk, J.; and Storksdieck, M. 2005. Using the contextual model of learning to understand visitor learning from a sci ence center exhibition. *Science education*, 89(5): 744–778.

[14] Council, N. R.; et al. 2009. *Learning science in infor mal environments: People, places, and pursuits*. National Academies Press.

[15] Stocklmayer, S. M., Rennie, L. J., & Gilbert, J. K. (2010). The roles of the formal and informal sectors in the provision of effective science education. *Studies in science education*, *46*(1), 1-44.

[16] Russell, J. L.; Knutson, K.; and Crowley, K. 2013. Infor mal learning organizations as part of an educational ecol ogy: Lessons from collaboration across the formal-informal divide. *Journal of Educational Change*, 14(3): 259–281.

[17] Long, D., Jacob, M., & Magerko, B. (2019, June). Designing co-creative AI for public spaces. In *Proceedings of the 2019 Conference on Creativity and Cognition* (pp. 271-284).

[18] Long, D., Padiyath, A., Teachey, A., & Magerko, B. (2021, June). The role of collaboration, creativity, and embodiment in AI learning experiences. In *Proceedings of the 13th Conference on Creativity and Cognition* (pp. 1-10).

[19] Dorph, R.; Cannady, M. A.; and Schunn, C. D. 2022. What drives visitor engagement in exhibits? The interaction be tween visitor activation profiles and exhibit features. *Cura tor: The Museum Journal*, 65(2): 399–416.

[20] Moore, R. L.; Jiang, S.; and Abramowitz, B. 2022. What would the matrix do?: a systematic review of K-12 AI learning contexts and learner-interface interactions. *Journal of Research on Technology in Education*, 1–14.

[21] Greeno, J. G. (1998). The situativity of knowing, learning, and research. American psychologist, 53(1), 5.

[22] Greenwald, Krakowski, Hurt, Grindstaff, and Wang. 2024. “It’s like I’m the AI”: Youth Sensemaking About AI through Metacognitive Embodiment. In ACM Interaction Design and Children (IDC) Conference. ACM, New York, NY, USA.

[23] Yin, R. K. (2009). Case study research: Design and methods (Vol. 5). sage.

[24] Stake, R. E. (2013). Multiple case study analysis. Guilford Press.