Co-Designing AI Literacy Exhibits for Informal Learning Spaces

DURI LONG

Expressive Machinery Lab, Georgia Institute of Technology, Atlanta, GA, duri@gatech.edu

TAKERIA BLUNT

Expressive Machinery Lab, Georgia Institute of Technology, Atlanta, GA, tblunt3@gatech.edu

BRIAN MAGERKO

Expressive Machinery Lab, Georgia Institute of Technology, Atlanta, GA, magerko@gatech.edu

AI is becoming increasingly integrated in common technologies, which suggests that learning experiences for audiences seeking a "casual" understanding of AI—i.e. understanding how a search engine works, not necessarily understanding how to program one—is an increasingly important design space. Informal learning spaces like museums are particularly well-suited for such public science communication efforts, but there is little research investigating how to design AI learning experiences for these spaces. This paper explores how to design museum experiences that communicate key concepts about AI, using *collaboration*, *creativity*, and *embodiment* as inspirations for design. We present the design of five low-fidelity AI literacy exhibit prototypes and results from a thematic analysis of participant interactions during a co-design workshop in which family groups interacted with the prototypes and designed exhibits of their own. Our findings suggest new topics and design considerations for AI-related exhibits and directions for future research.

CCS CONCEPTS • Applied computing~Education~Interactive learning environments • Social and professional topics~Professional topics~Computing education • Computing methodologies~Artificial intelligence

Additional Keywords and Phrases: **AI literacy, embodiment, collaboration, creativity, co-creativity, AI education, informal learning, co-design, participatory design**

ACM Reference Format:

Duri Long, Takeria Blunt, and Brian Magerko. 2021. Co-Designing AI Literacy Exhibits for Informal Learning Spaces. In CSCW '21: ACM Conference on Computer Supported Cooperative Work and Social Computing, October 23–27, 2021, held virtually. ACM, New York, NY, USA. NOTE: This block will be automatically generated when manuscripts are processed after acceptance.

1 INTRODUCTION

Artificial intelligence (AI) is becoming increasingly integrated into our day-to-day experiences—not only for individuals who choose to purchase AI devices like smart speakers or self-driving cars, but also for people who interact with AI in service of other goals—to unlock their phone, interact with friends on social media, search for something on Google, or complete their job at work. Recent research has shown that there are numerous public concerns related to the presence of AI in commonly used applications—including personal data privacy breaches [82], bias and discrimination in algorithms [15], and issues with viral misinformation [3,72]. Public education efforts have historically played an important role in fostering a scientifically and technologically literate populace who can critically engage with technology and advocate for policy and design choices that support their needs and values [84].

Most AI education efforts to-date have focused on teaching individuals—either university students in AI-related degree programs or, more recently, K-12 students in introductory computing lessons—how to program AI [25,89]. In this paper, we focus on designing for people (both adults and children) who may not necessarily be interested in learning about how to program AI—but who could benefit from a high-level understanding of AI in their day-to-day interactions with technologies. We have previously defined this 'high-level understanding' as *AI literacy*, or "a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace" [60]. AI literacy is a distinct set of skills, although it can be informed by other related literacies, such as scientific or computational literacy, and has some overlap with data literacy, which is highly relevant to the AI subfield of machine learning [60].

Informal learning spaces—like museums—have long served an important role in public science communication efforts. Science and technology education in public spaces like museums can aid in content knowledge gains [4,5], interest development [42], and improved understanding of connections between science/technology and society and culture [84]. Although there have been some AI-focused museum and art exhibits (especially in the past few years), there are still many aspects of AI that have yet to be explored in these contexts, and there are few existing projects that draw on evidence-based research and/or have resulted in findings that have been disseminated to the research community.

Findings from both research on designing museum exhibits and HCI research on developing technology-based installations for public spaces suggest that museums/public spaces pose a number of design constraints and opportunities that are distinct from formal learning environments [4,29,40,45,48,83]. Most contemporary science museums are designed to facilitate open-ended, visitor-led learning experiences—as opposed to teacher-led classroom lessons [29]. Interaction times may be brief and most visitors travel in groups, whose attention can easily shift from exhibit to exhibit [40]. Museums focus on outcomes such as fostering interest development and facilitating introductory exploratory learning experiences (often with a variety of possible learning outcomes depending on differing visitor goals) [29,42]. This contrasts with formal learning environments, which more often focus on teaching a structured curriculum.

Previous research has indicated that certain design features can aid in quickly engaging visitors in active learning experiences. Incorporating *embodied interaction* can make exhibits easily understandable to visitors with little prior knowledge in addition to facilitating an engaging experience [45–47,76,83]. Open-ended, *creative* exhibits encourage prolonged engagement in addition to facilitating visitor-led learning experiences that can lead to more personally relevant meaning-making [11,29,49]. Supporting *collaboration* and social interaction between multiple group members creates a social learning environment in which visitors can learn from and motivate each other [23,49,96]. Taken together, *embodied interaction*, *creativity*, and *collaboration* may lead to more engaging and effective museum learning experiences. Early research on AI education as well as prior work in computer science education indicates that embodied interaction, creativity, and collaboration may also be particularly well-suited to promoting learning and interest development in computing related fields [1,14,23,25,39,65,85,95,101].

This paper investigates the following research question: *How can we design embodied, co-creative learning experiences for museums that foster interest in and improve public understanding surrounding AI?* We used existing guidelines for designing AI literacy learning experiences [60,88] as a foundation for the design of five AI-related exhibit concepts and low-fidelity design prototypes. We present these prototypes as an initial design exploration into how to promote AI literacy in informal learning spaces. We then present the results from a study in which we engaged stakeholders (family group museum visitors and museum staff at a large science and technology museum

in the midwestern United States) in co-design activities in order to: 1) better understand our target audience's prior knowledge about AI; 2) gather feedback on and insights into participant interaction with the exhibit prototypes we developed; and 3) generate new design directions for AI-related museum exhibits. The results from this study are presented in the form of common themes that emerged from a thematic analysis [69] of the data.

This paper's contribution is two-fold: 1) a set of low-fidelity AI literacy exhibit prototypes, with evidence-based directions for future design research and 2) an emerging understanding of visitor preconceptions about AI, important design features for AI literacy exhibits, and learning talk that occurs at AI literacy exhibits. Our work is in conversation with prior work in the CSCW community that investigates the role that collaboration can play in the design of public exhibitions and informal learning experiences (e.g. [43,87]) as well as research investigating shared perceptions and narratives related to AI (e.g. [64]). We also share publicly available versions of our exhibit designs and workshop materials for researchers and designers in the field interested in replicating or building on this work at this link.

2 RELATED WORK

2.1 AI Education

AI education research for K-12 and other non-expert audiences is still in its early stages, but there has been a growing interest in this space, particularly in the past several years. Touretzky et al. outlined five "big ideas" of AI, setting the agenda for research on K-12 AI education—1) "Computers perceive the world using sensors"; 2) "Agents maintain models/representations of the world and use them for reasoning"; 3) "Computers can learn from data"; 4) "Making agents interact with humans is a substantial challenge for AI developers"; and 5) "AI applications can impact society in both positive and negative ways" [88]. Our review of existing AI education research built on Touretzky et al.'s work to synthesize a set of AI literacy competencies (e.g. "Understand what a knowledge representation is and describe some examples of knowledge representations") and design considerations (e.g. "Consider including graphical visualizations, simulations, explanations, of agent decision-making processes, or interactive demonstrations in order to aid in learners' understanding of AI") [60]. We use Touretzky et al.'s five big ideas and our previously defined AI literacy competencies/design considerations as guiding frameworks for our design research in this paper.

Several projects have aimed to make learning about AI accessible to broader audiences by providing toolkits that enable learners to "tinker"—or learn through active experimentation—with existing AI devices such as Siri, Alexa, Google Home, Cosmo, Roomba, and other robots [22,25,79,88]. Another genre of approaches to teaching AI involves interactive demonstrations of how certain AI technologies work "behind the scenes" (e.g. [36] [37]). AI education for more advanced learners can involve creating AI programs or building machine learning models from scratch. Researchers are beginning to explore how to design interactive tools and learning environments that can scaffold these more advanced activities [85,101]. Finally, some projects explore how to teach AI ethics concepts in connection with learners' prior knowledge of common technologies that use AI, like YouTube [21]. Our design research draws on findings from this prior work while also investigating new approaches to designing AI literacy interventions specifically for museum spaces.

There have been only a few public-facing exhibits focused on AI, likely due to the novelty of the field, the expense/fragility of many AI devices, and the fact that the 'inner workings' of most AI are not easily interactive or observable. Most often, AI appears in science centers in the form of displays or demonstrations of robotic artifacts—

for example, the Museum of Science and Industry, Chicago's *Robot Revolution* exhibit featured numerous historical robotic artifacts and demonstrations of robots performing feats of intelligence (e.g. reciting Shakespeare) [67]. However, robotics-focused exhibits often neglect other areas of AI (i.e. cognitive systems, machine learning), and facilitating visitor-led interactive learning experiences can be difficult due to the expensive, easily breakable, and sometimes physically dangerous nature of robotic devices.

Recent work is exploring how to make robotics-related exhibits more interactive and interdisciplinary. *Cubelets* facilitate a safer, more durable, exploratory learning experience with robotics by allowing learners to make simple robots by connecting modular magnetic sensor and actuator cubes [17]. A recent project developed for the Carnegie Science Center demonstrates how a self-driving car might work using the commercially available Cozmo robot and real-time visuals of the robot's world map and image recognition [74]. This project is unique in that it communicates information not just about robots, but also other AI topics like knowledge representations and computer vision.

Installations in art spaces have also explored how to communicate key ideas about AI. Several art installations have touched on ethical topics related to AI, aiming to engage non-technical audiences in critical discussions about AI [18,41]. Other art spaces have curated multiple AI-related installations—including artifacts, artwork, and interactive demonstrations—into overarching exhibits that explore more holistic representations of AI [9,86]. However, there are still many aspects of AI that have yet to be explored in a museum context.

2.2 Co-Design and AI

Participatory processes (i.e. processes that actively involve stakeholders in the design process) have been discussed as being critical to the design of artificially intelligent systems due to cultural, social, political and ethical implications behind AI technologies [57]. As the applications of AI span many disciplines, contexts, and uses, practitioners in education, healthcare, design, computer science, engineering and beyond are exploring methodologies for better involving humans in the co-construction of these systems. While not focused specifically on education, projects delve into using participatory methods such as contextual inquiry in combination with other methods to develop explainable, equitable, and human-centered AI [77,81,93]. Past projects have engaged in low-tech prototyping with youth groups to develop an AI librarian [81], explored how to engage cognitively impaired communities in the design of social robotics [77], diversified human involvement in the training of AI models used in critical decision making that disproportionately impacts marginalized groups of people [30], and engaged users in co-creating human-centered explainable AI [93]. We drew on methods from this prior work in our approach to the co-design study presented in this paper, which uses techniques such as low-tech prototyping, soliciting stakeholder input and feedback on design concepts, and co-creating novel designs with stakeholders.

Although not specifically focused on AI, research on co-designing technology with family groups is also relevant. Gutwill & Allen studied family groups' scientific inquiry practices by engaging families in interactive game activities. In designing these games, they drew on numerous design principles including building on learners' prior knowledge, supporting collaboration, explicitly identifying a manageable number of skills for each activity, supporting metacognition/reflection, and supporting a flexible, intrinsically motivating learning agenda that can adapt to visitor's personal needs [38]. Others have studied how to support co-design amongst family groups and have found that family-based co-design is shaped and impacted by parent-child dynamics, parenting styles, and other relational nuances that require further exploration [97]. Design considerations such as supporting creative expression and individual as well as collaborative design ideas have been found to support intergenerational co-design of technology [92]. Dindler et al.'s work additionally suggests that incorporating children's existing "funds of

knowledge" can contribute to exhibit ideation/design processes [20]. We draw on some of the design principles from these studies in the design of our workshop and also in interpreting our results in the Discussion (section 8).

2.3 Designing for Museums

Research on museum design has also informed the work presented in this paper. Informal learning spaces such as museums are heavily influenced by the theory of construct*iv*ism, which argues that knowledge is constructed by the learner via active engagement with the environment [11]. The idea that knowledge is actively constructed by the learner was originally suggested by Piaget [7]. Vygotsky, Bruner, and other learning scientists have expanded on this definition of constructivism to suggest that all knowledge is *socially* constructed—that is, learning is an active, social process in which learners construct knowledge through interaction with others and their environments [7,54,90]. Viewed through the lens of social constructivism, *learning* is an active process of meaning-making that involves connections to both social and personal identity. This notion of learning has inspired and guided the design of many museums and science centers that frame themselves as *free-choice* learning environments, or environments in which learner agency is prioritized and exhibits are frequently designed to facilitate embodied social interaction [29]. Many projects over the years in the HCI and CSCW communities have explored how to design collaborative installations involving technology for museum spaces (c.f. [48]). Of particular relevance to our work is research that has explored how to design exhibits for CS learning [44,45,61,62,70], as well as recent work that has begun to identify some of the unique design challenges and opportunities involved in introducing AI in public spaces [59]. Also relevant is research on studying families' sense-making processes and learning talk surrounding museum exhibits [38,76,99].

3 AI LITERACY EXHIBIT PROTOTYPES

3.1 Methodological Approach

We use *design research*—or "[research] in which design practice is brought to bear on situations chosen for their topical and theoretical potential…[where] the output takes the form, primarily, of artefacts and systems, sometimes with associated accounts of how these are used in field tests"—as the primary methodological approach for exploring how to design museum exhibits to communicate AI literacy competencies [33]. Zimmerman et al. suggest that design research in HCI should have a well-documented *process*, constitute a significant *invention* with respect to prior literature, be *relevant* to the context, and be *extensible* (i.e. provide a foundation for future research). We aim to meet these criteria by documenting our design process and motivation in detail, situating our designs with respect to prior research on AI education, museum education, and social AI literacy more broadly, and synthesizing design themes that can be used to inspire and guide future work in the field [100].

We used three guiding principles as inspiration for our exhibit design artifacts—*embodied interaction*, *collaboration*, and *creativity*. We focused on these principles because they have been shown to encourage learning and are particularly well-suited to fostering engagement in a museum environment. Not all of the exhibits we developed incorporate all design features (embodiment, collaboration, and creativity), but each exhibit incorporates at least one. Each of the principles is defined and described in more detail below.

Embodied Interaction: Dourish defines embodied interaction as "the creation, manipulation, and sharing of meaning through engaged interaction with artifacts'" [24]. *Embodied interaction* [24,51] has been shown to promote learning in several domains [1,8,25,28,85] and is an engaging, easy-to-understand way to interface with

novel systems in museums [4,76,83]. In this paper, we explore designs that utilize full-body interaction, tangible user interfaces (TUIs), and embodied spatial metaphors.

Collaboration: There are a variety of different definitions of collaboration in the literature involving differing degrees of shared attention and social interaction. We draw on a framework developed by Ludvigsen [63] that identifies four levels of interaction that occur with social interfaces, based on Goffman's sociological research [34]. *Distributed attention* refers to a group of people who are sharing the same space, but who are not focused on any shared interface or activity. *Shared focus* is when a group of people is focusing on a single activity/interface within a space. *Dialogue* occurs when people engage in a shared activity, and *collective action* occurs when people work together to achieve a shared goal. In this paper, I define collaboration as encompassing both *dialogue* and *collective action* and explore designs that can facilitate these types of interactions. Collaboration between peers and family members both has been shown to encourage learning and interest development in a variety of learning environments (e.g. [23,95]) and is especially important in museum settings since most visitors come in groups [40].

Creativity: Boden defines creativity as ~the ability to generate novel and valuable ideas" [12]. Ideas can be novel to the person generating them (*P-creative*) or historically novel (i.e. never occurred in history before; *H-creative*). Creativity can take on a variety of different forms, including combining familiar ideas in novel ways (*combinational creativity*), exploring a conceptual space (*exploratory creativity*), or transforming a conceptual space (*transformational creativity*). Open-ended activities that allow for creativity and personal expression have been shown to increase interest formation in computing [14,39,65,75]. In addition, museum exhibits that allow for creativity have led to lengthier interaction times and increased learning gains [49]. In this paper, we explore designs that encourage learners to generate P-creative ideas by expressing themselves through activities like dance, drawing, or generating novel artifacts and/or combinations of ideas.

We additionally drew on findings from prior work on AI literacy and AI education for K-12 audiences during our design process. We use AI literacy competencies developed in our prior work [60] and the "big ideas" of AI from [88] to specify learning goals for our exhibit designs, and we used the design considerations from our prior work [60] as addition inspirations for design. The next section will detail which competencies, design considerations, and "big ideas" are incorporated in each exhibit prototype.

3.2 Exhibit Design

Table 1: Summary of competencies and design considerations in each exhibit prototype.

Exhibit Name	Competencies	Design Considerations
Magic Mirror	understanding sensors	low barrier to entry; embodied interaction
Sensor Wall	understanding sensors	low barrier to entry; embodied interaction
Neural Net	steps and practices of machine learning; ways of representing knowledge; how agents make decisions	embodied interaction; opportunities to program or teach AI; explainable algorithms
Semantic Network	ways of representing knowledge; how agents make decisions	embodied interaction; opportunities to program or teach AI; explainable algorithms; incorporating learner interests; opportunities for collaboration
LuminAI	steps and practices of machine learning; ways of representing knowledge; how agents make decisions; computers learn from data; different types of intelligence	embodied interaction; explainable algorithms; opportunities for collaboration; opportunities to program or teach AI; incorporating learner interests; engaging with lesser-known forms of AI

We constructed low-fidelity prototypes of five different exhibits: *Magic Mirror*, *Sensor Wall*, *Neural Net*, *Semantic Network*, and *LuminAI*. These exhibit concepts were originally brainstormed as part of a *design workbook*, or a "collection...of design proposals and other materials drawn together during projects to investigate options for design" [32]. We selected these five exhibits from a set of 24 different ideas in the workbook based on their relative novelty in comparison to existing work, the variety of ways they employed embodied interaction, collaboration, and creativity, and their ability to collectively represent competencies spanning all five of the "big ideas" of AI and incorporate a variety of the AI literacy competencies and design considerations outlined in our prior work [60]. We wanted to get a sense of whether families could grasp complex concepts like aspects of how neural networks work during a short interaction, but also wanted to incorporate themes from the other "big ideas" of AI, such as understanding sensors. Each exhibit concept and its prototype is described in more detail below, including a list of the "big ideas" and/or AI literacy competencies/design considerations that it incorporates and the ways in which it incorporates embodied interaction, collaboration, and creativity. This information is summarized in Table 1.

3.2.1 Magic Mirror

Figure 1: Magic Mirror. From left to right: Exhibit concept design sketch, prototype setup, example cyborg artifact.
Image on left © Jonathan Moon & Duri Long

Magic Mirror is an exhibit in which learners are encouraged to use their knowledge of their own body (i.e. "body sintonicity" [2]) to learn about the ways in which a computer senses and acts on the world. In the conceptual exhibit design sketch, learners would look into a mirror and see computational sensors and actuators that would perform similar functions to their organs projected onto their body. For instance, a visitor may see capacitive touch sensors where their hands would be, sound sensors on their ears, and a speaker on their mouth (). The low-fidelity paper prototype that we shared at the workshop was an activity in which participants were given an outline of their face or body as well as a deck of sensor/actuator cards. Each card had an image of a sensor or actuator on the front and a description of the device on the back. The participants were prompted to consider several questions: 1) Can you make a "cyborg" (a human-machine hybrid) using human parts, computer sensors, and actuators?; 2) If you could replace one body part with a computer sensor/actuator, what would it be?; 3) What would you never replace?; 4) How are sensors similar or different to your five senses?; and 5) If you had to build a human body using sensors and actuators, what would you put where?. As part of the first question, the participants created "cyborgs" with

human and machine body parts by gluing cards onto face or body outlines. An example of a finished artifact is shown in . The low-fidelity prototype was conceptualized as a way to gather feedback on aspects of the concept in the design sketch without building complex image recognition software; however, as we expand upon in Results (section 7) and Discussion (section 8), the low-fidelity prototype may have been more successful at supporting collaboration/creativity than the conceptual design sketch.

This exhibit concept aims to communicate AI-related competencies such as *understanding different types of sensors/actuators* and their capabilities [60,88]. This exhibit is *embodied* in that it encourages learners to draw on knowledge of their own body to help make sense of how sensors and actuators are similar to or different from their own sensory organs. The low-fidelity prototype is open-ended and *creative* in that it allows learners to consider how to creatively leverage human and machine capabilities to make an imaginative cyborg (the conceptual exhibit in the design sketch is less open-ended). The prototype allowed for family members to *collaboratively* make a cyborg together (i.e. *collective action* as they work towards a shared goal), and both the prototype and the design sketch are designed to encourage collaborative discussion of sensors and their capabilities (i.e. *dialogue*) [63]. This exhibit (particularly the conceptual exhibit sketch) also has a *low barrier of entry*—even very young children can walk up to the mirror and see the sensors/actuators projected onto their body [60]. However, the exhibit also allows for *multiple levels of engagement* including fostering in-depth discussion of different sensors/actuators, their relationship to human body parts, and the practical and philosophical implications of cyborgs. This allows for collaborative engagement between family members of all ages.

3.2.2 Sensor Wall

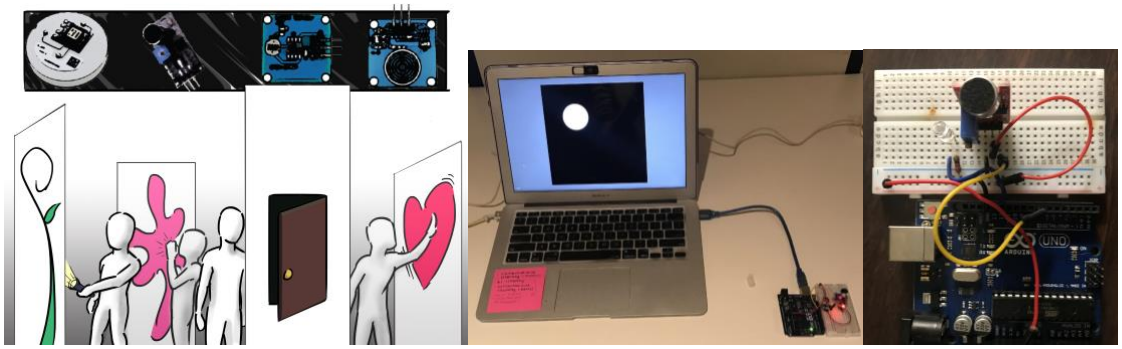

Figure 2: Sensor Wall. From left to right: Exhibit concept design sketch, prototype setup, zoom in on sound sensor.
Image on left © Jonathan Moon & Duri Long

Sensor Wall is an exhibit concept that uses embodied interaction to engage visitors in learning about different sensors and their capabilities. Learners can interact with simple, responsive AI projections on walls. Visitors can shine a flashlight on a light sensor to make a projected plant grow, shout at a sound sensor to make a blob explode, approach a proximity sensor to make a projected door open, and place their hand on a touch sensor to make a heart glow (). For the workshop, we provided a simple prototype for participants to interact with where they could make a noise into a microphone and in response, a LED would light up and a circle would blink on the computer screen ().

This exhibit concept aims to communicate AI-related competencies such as *understanding different types of sensors* and their capabilities [60,88]. The exhibit is *embodied* in that it encourages learners to engage in physical exploration using their bodies and provided tools to make sense of sensors and how they react to environmental changes. It draws on a tradition of "social immersive media" that involves visceral interaction—that is, interaction that "is first understood by the body and later understood rationally" [83]. This exhibit also builds on prior work on exhibit design where visitors can explore scientific phenomena using their bodies and embedded tools [49]. The exhibit also allows for multiple users to interact with the sensors together, which has the potential to facilitate collaboration as learners engage in *collective action* to elicit certain responses from the sensors or just engage in verbal *dialogue* about the sensors and how they work [63]. *Sensor Wall* also facilitates a *low barrier of entry* for younger visitors, allowing for collaboration among group members of all ages [60].

3.2.3 Neural Net

Figure 3: Neural Network. From left to right: Exhibit concept design sketch, prototype setup, interaction with prototype. Image on left © Jonathan Moon & Duri Long

Neural Net is an exhibit concept in which visitors can interact with a physical representation of a neural network[1] (Figure 3). Each node in the network represents a feature, and putting physical weight on that feature by pressing it corresponds to training the network on more data with that feature, thereby placing an emphasis on that feature in the network . In the example shown in , the neural network is a bird classifier. Placing more weight on a feature like "ability to fly" would mean that the network is trained on mostly examples of flying birds, causing birds like kiwis to be misclassified as "not a bird". Visitors can explore different inputs and weights to see what outputs result. We shared a low-fidelity physical prototype of the exhibit combined with an interactive demonstration at the workshop (). The exhibit was prototyped using a wooden box with elastic strings stretched across it. At each node, a circle was glued with an image of a particular bird feature (e.g. wings, beak). There were cards on the table with images of different birds as well as animals that might be mistaken for birds (e.g. turtle, platypus). We also included some atypical birds that might be misclassified (e.g. penguin, ostrich). There were also two index cards provided—one said "It's a bird!" and the other said "It's not a bird!" Participants were given an interactive demonstration of what the system might do in certain scenarios (e.g. classify a robin as a bird but a penguin as not a bird if emphasis was placed on the flying feature). Participants were then encouraged to look through the deck of animal cards and

1 The version of this prototype presented at the workshop used language related to neural networks. However, focusing on human-described features when describing neural networks—which learn directly from input image data and therefore may not describe birds in the same way humans do—could potentially be misleading [27]. In iterations built after this workshop, we renamed this prototype *Creature Features* and use a feature-based learning algorithm rather than a neural network to classify birds.

think about how they could get each to be classified or not classified as a bird by adjusting the weights in the network.

This exhibit concept aims to communicate AI-related competencies such as the *steps and practices of machine learning*, *ways of representing knowledge*, and *how agents make decisions*. This exhibit incorporates *embodiment* in two different ways. First, it utilizes a tangible interface, concretizing an otherwise abstract and "black-box" algorithm (embodied interfaces have been used successfully in other fields to facilitate learning of abstract concepts—e.g. [1,16]). In addition, the exhibit employs a "weight" metaphor [53], drawing a connection between weighted and biased training data and the embodied notion of weight. This is intended to help users draw on their "body sintonicity" (i.e. knowledge of their body and how it interfaces with the world) when making sense of the network [2]. The exhibit is also *collaborative*—learners are encouraged to work together to decide which features to put weight on (*collective action*) and engage in verbal *dialogue* about the possible outcomes of the network [63]. The physical design is large enough to allow multiple users to place weight at once, and it is intended to encourage group members to debate/discuss where they should place the weights to best optimize the network. This type of constructive debate—particularly in a family group collaboration—can increase the likelihood of conceptual change [23]. *Neural Net* also incorporates design considerations such as *providing opportunities to program or teach AI*, and *creating explainable algorithms* [60].

3.2.4 Semantic Network

Figure 4: Semantic Network. From left to right: Exhibit concept design sketch, prototype setup, completed artifact.
Image on left © Jonathan Moon & Duri Long

Semantic Network is an exhibit concept in which visitors can build semantic networks that describe concepts of interest to them (e.g. their family, relationships between things in the world) using strings, pins, and a corkboard (Figure 4). Different string colors indicate different relationships (e.g. is, has, likes, dislikes). Visitors can take a picture of their creation and it will be parsed by a computer. The visitor can then ask the computer questions about the semantic network that they created (e.g. Q: "What does a cat have?" A: "A cat has fur"). By building a network and asking the AI questions about it, participants can explore/discuss ideas including the limitations of knowledge representations, the differences between human and computer intelligence, and what it means to "know" or "understand" a concept. In the workshop, participants were given a set of circle cards with entity/concept names on them (e.g. mom, dog, piano) and a set of colored lines. They were also given a key—each colored line corresponded to a different type of relationship (e.g. is-a, has-a, likes, dislikes). Participants were prompted to

create their own semantic network to teach a hypothetical computer. Participants then glued down the appropriate concepts and lines according to the key.

This exhibit concept aims to communicate AI-related competencies such as *ways of representing knowledge* and *how agents learn and make decisions*. This exhibit incorporates *embodiment* in the form of a tangible interface for constructing a network. Physically building the network in space aids in concretizing the abstract and otherwise "invisible" nature of knowledge representations. The exhibit is designed to be *collaborative*—multiple group members are able to easily work together to build a single network (i.e. *collective action* [63]). There are multiple points of entry—little kids can think about relationships they may currently be learning (e.g. cat has whiskers). Older kids and adults can expand the network's complexity and explore its limitations. The exhibit is *creative* in that it is open-ended, and users can explore how to create personally meaningful networks depending on their interests. *Semantic Network* also incorporates design considerations such as *opportunities to program or teach AI*, *creating explainable algorithms*, and *leveraging learner interests* [60].

3.2.5 LuminAI + MoViz

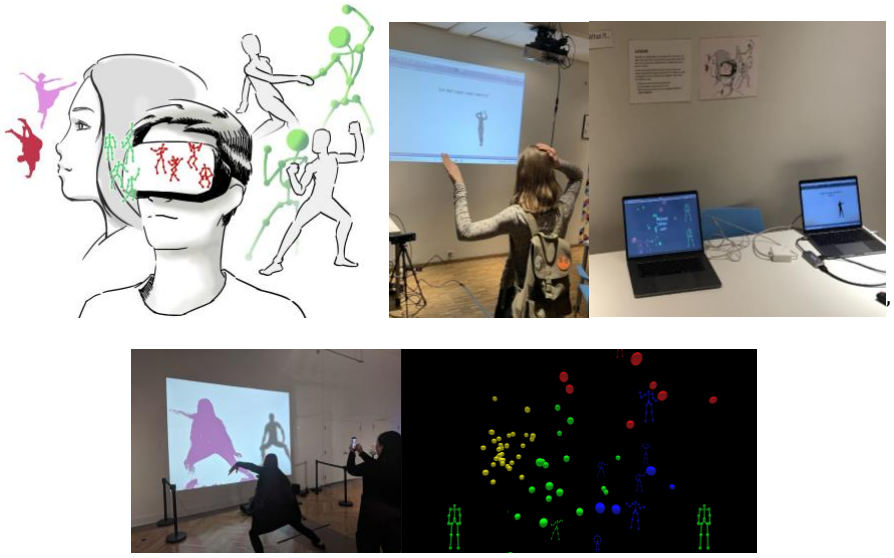

Figure 5: *LuminAI*. From left to right: Exhibit concept design sketch, prototype setup at workshop, close-up of *LuminAI* dance interface, close-up of *MoViz* visualization software. Image on top left © Jonathan Moon & Duri Long.

The *LuminAI* exhibit concept builds on an existing AI installation in which participants can improvise movement together with an AI dance partner that is projected onto a screen () [59]. A "shadow" of the user's body (generated from a motion capture sensor) is projected on the screen next to an AI dance partner, which senses participant movements via a motion capture sensor and responds with movements in its memory that it deems to be similar. This exhibit concept proposes expanding *LuminAI* into an educational exhibit where participants can move between multiple interaction stations to create a personally customized AI dance partner and learn about the AI dancer's gesture memory. Participants would be able to switch between datasets of different genres of dance moves (e.g. hip-hop, ballet) in order to change the AI agent's knowledge of dance. Participants can put on a virtual reality (VR) headset and use a tool called *MoViz* to explore a visualisation of the gestures in a chosen dataset, which are clustered

(grouped) based on their similarity (e.g. a right hand wave and left hand wave would be clustered together) [56]. Visitors can then dance with an AI dancer (projected onto a screen or wall) that is trained on the dataset they selected (). *LuminAI* was at a more developed stage prior to the workshop than the rest of the prototypes. We allowed participants to interact with the original system by dancing with an AI dancer projected on the wall, and then had them interact with *MoViz* using a laptop interface (). At the time of the workshop, *LuminAI* and *MoViz* were not communicating—that is, *MoViz* was displaying a different dataset than the one *LuminAI* was using to dance.

The *LuminAI* exhibit concept touches on competencies such as *understanding the steps and practices of machine learning*, *ways of representing knowledge*, *how agents make decisions*, *how computers learn from data*, and *understanding different types of intelligence*. *LuminAI* incorporates *embodiment* in two different ways. First, the exhibit utilizes an embodied interface in which learners use their bodies to dance with the AI. Second, the *MoViz* visualization of gesture clustering is an embodied representation of the way in which the agent groups similar gestures in memory. Learners can explore gestures and their clusterings in relation to each other in space, concretizing and visualizing learned relationships between data points that would otherwise be invisible to the user. The exhibit can be highly *collaborative* when a setting is enabled to allow multiple user "shadows" to appear on the screen next to the AI agent—in the past we have observed a variety of emergent social dance interactions, including physical *dialogue* as visitors engage in dance together and *collective action* as visitors work together to get the agent to respond in a certain way [58,63]. However, the prototype we presented at the workshop only allowed for single user interactions due to space constraints. While not strictly collaboration according to the definition presented at the beginning of this paper, *LuminAI* also learns from users and repeats their dance movements back to later participants, allowing for a shared dialogue across time with others that have danced with the exhibit throughout the day. Finally, the exhibit is *creative* not only in that it is situated in a creative domain that allows for personal expression, but also in that it allows learners to co-create together with and customize an AI dance partner. *LuminAI* also incorporates additional design considerations such as *creating explainable algorithms*, *providing opportunities to program or teach AI*, *incorporating learner interests*, and *engaging with lesser-known forms of AI* [60].

4 METHODOLOGY

4.1 Site Location and Target Audience

The site location for our research was the Museum of Science and Industry, Chicago a large science and technology museum in the midwestern United States.. In 2018, over 1.3 million guests visited the museum [73]. We chose this museum as a site location due to its large audience, ongoing work by their development team on a series of AI-related exhibits, and an existing research partnership. Since this is an exploratory study in an emerging field of research, we focused on recruiting a relatively small number of participants that we could spend a long time with in order to gather in-depth, qualitative data on family groups' needs, goals, and prior knowledge in relationship to AI and AI-related learning experiences. Our target audience was family groups visiting the site location. We focused on recruiting families with children over the age of six, since most children do not develop theory of mind (i.e. the ability to "explain and predict other people's behavior by attributing to them independent mental states, such as beliefs or desires" [31]) until they are 3-5 years old, making it difficult to grasp AI as a separate entity [94]. However, we recognized that some families may have younger children who they wanted to bring along to the museum, so we provided some blocks and craft materials (which were also used for the design activity) in case very young children grew bored or distracted.

We focused on studying family groups for several reasons. First, 75% of museum visitors come in groups [40], and providing opportunities for collaborative interactions was one of our guiding design principles. In addition, we hoped that by focusing on family groups, we could explore how to expand AI education for both children and their parents, as adult education is an important and often overlooked component of public science education [66] and very little research to date has focused on AI education for adults without backgrounds in computing or data science. We aimed to design exhibits with multiple levels of entry, meaning that they would be easy to understand for younger children, but would offer opportunities for deeper engagement for older children and adults. We are interested in fostering learning/interest outcomes for adults and children alike (i.e. we are not exclusively focused on K-12 educational outcomes but rather the collaborative learning outcomes for all members of the family group).

4.2 Study Setup

We conducted two co-design workshops which were each two hours long and were held in the morning on different days in a small classroom workspace in the museum. Two researcher-facilitators collaboratively ran each workshop. Participants were recruited using a museum mailing list of visitors who had previously expressed interest in engaging in studies. We also sent recruitment emails to museum staff and researchers in order to involve on-the-ground expertise in our design sessions. Each participating family was compensated with free admission to the museum for up to four family members, free parking, and one free "add-on" experience (an estimated value of ~$160). The workshop was conducted in the morning so that participants would have plenty of time to explore the museum after the workshop.

The workshop held on Day 1 had just one family in attendance due to a last minute cancellation. The Day 2 workshop consisted of four families and one museum team member. A number of additional museum staff members chose to attend an informal feedback session after the Day 1 workshop due to time constraints (the feedback from this session is not reported in this paper due to its informal nature). Participants worked with their family groups throughout the workshop, and we paired a museum team member with one of the families on Day 2. Across the two days, we had a total of 20 participants—nine adults and 11 children. Ten of these children were ages six and up. Each group is numbered and described in more detail in Results (section 7). We did not collect any demographic data on participants beyond group size and approximate ages of children, which was shared when participants signed up for the study for planning purposes.

4.3 Data Collection

We collected a variety of data for analysis in different mediums during the workshop (as is common in generative design research [78]). We photographed/scanned artifacts generated by participants—including worksheet packets, design sketches, written prototype feedback, and exhibit prototype artifacts. One of the researcher-facilitators present at the workshop took field notes during and after the workshop. We also recorded audio of participant interactions during the workshop to corroborate our notes and provide additional detail on participants' thought processes and commentary. Audio recorders were placed on participant group tables and later moved to prototype stations (see Workshop Summary (section 5) for details on the workshop structure). Most of the participants' verbal dialogue was captured in audio recordings with a few exceptions due to inaudible/overlapping dialogue and the difficulty of moving the limited number of audio recorders between tables in a timely manner. Data from field notes and the collected design artifacts provide redundancy and supplement the data where recorded dialogue was missing or inaudible.

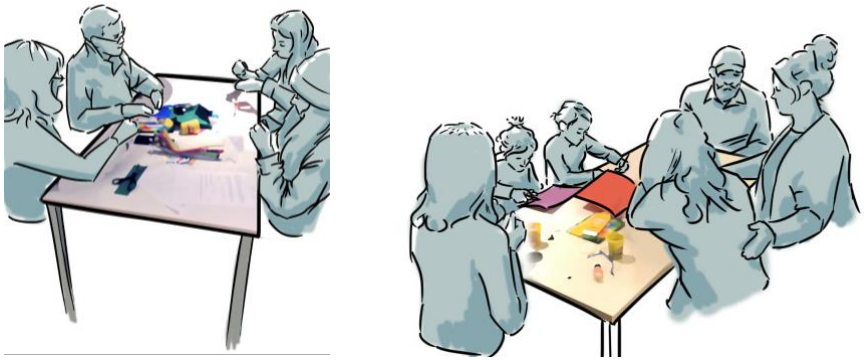

Figure 6: Sketches of workshop participant groups. © Jonathan Moon & Duri Long.

5 WORKSHOP SUMMARY

In this section, we describe each phase of the workshop in detail, including our motivation for the activities we included in each phase. We have made all of the materials used in the introductory and design activities as well as some of the low-fidelity prototype materials available online at this link. We hope that these can be of use to researchers and designers interested in conducting similar workshops and/or educators interested in using any of these materials in learning environments.

Research involving co-design can vary widely in terms of the level of stakeholder involvement, ranging from involving stakeholders in generating the research questions to engaging stakeholders in later stages of prototyping. We purposefully incorporated varying degrees of stakeholder involvement in our workshop. We came into the workshop with existing research questions and several flexible design ideas in the form of low-fidelity prototypes, but we also wanted participants to generate exhibit design ideas of their own. We originally intended to show participants our exhibit prototypes after they had generated ideas of their own (so as not to influence their ideas), but after pilot testing the workshop activities with several different groups, it became clear that we needed to scaffold the design generation process to ensure that participants were able to build on a) some prior knowledge of AI and b) an understanding of design possibilities/constraints for museum exhibits. Situating the open-ended co-design activity at the end of the workshop ensured that participants had foundational knowledge of AI from both the worksheet activity and interaction with the prototypes and that they had some inspiration to draw from in regards to possibilities for exhibit design.

One of the reasons we chose to use the methods we did was because we had a limited amount of time—our partners at the Museum of Science and Industry advised us based on their prior experience that we would likely be unable to recruit participants for a workshop lasting longer than two hours (or spanning multiple days). This imposed some constraints on the types of activities we could incorporate—for example, iterating on participant-generated designs using techniques like layered elaboration [91] was not feasible within the timeframe. The activities we chose all had a low barrier of entry and required only minimal explanation, allowing participants to spend the majority of the workshop engaging in dialogue about AI and exhibit designs. We also aimed to incorporate activities that learners of all ages could participate in.

5.1 Consent

As participants arrived at the workshop, a researcher discussed the informed consent process with them and obtained written consent from adults and assent from minors. Written assent was obtained for children over the age of 11 and verbal assent was obtained for children under the age of 11, using age-appropriate language. Parents were also asked to sign a photo release form for the family. All participants were informed during the consent process that their conversations during the workshop would be audio-recorded and that they may be photographed.

5.2 Introductory Activity

After all participants had been consented, they were seated at a table with their group. The workshop began with an introductory activity. Participants worked together on a three page worksheet activity involving a card deck. All questions were answered as a group. All groups completed pages 1 and 2 of the worksheet activity. Group 2 arrived late and did not have time to finish page 3 of the activity. This entire activity took ~20 minutes for groups to complete.

The introductory activity served several purposes—it eased the groups into co-creation and ideation, provided us with some insight into their prior knowledge about AI, and equipped participants with foundational knowledge of AI that they could use in the later activities. The worksheet format allowed us to provide scaffolding in case participants did not know much about AI. This eliminated the need for an introductory presentation about AI and instead immediately engaged participants in collaborating and brainstorming together. The card deck was also intended to provide participants with many examples of AI and AI components to use later on during the design activity. The three worksheet pages are described in more detail below.

On the first page of the worksheet, we asked participants several questions about their prior experiences with AI and their impressions/feelings about AI. On the second page of the worksheet, participants were asked to look through a card deck (Figure 7) containing a number of examples of AI technologies and were prompted to select the ones they had interacted with before and glue them on the page. They were then prompted to circle items they were surprised by. This activity was intended to give participants an idea of what technology they interacted with actually used AI (helping resolve commonly held confusion about what AI is [60]). We also wanted to gauge which (if any) technologies were surprising to participants. On the third page of the worksheet, participants were asked to construct an imaginary AI technology by choosing three components: an input, an algorithm, and an output. They were given four different card decks to select from (Sensor and Dataset decks for the input component and Algorithm and Output decks for the latter two components, pictured in Figure 7). Participants walked through the activity step by step and learned about the roles of each component in the instructions as they went along. This activity was intended to familiarize participants with the basic components of AI so they would have an idea of how AI works before getting involved in the prototype testing and design activities. We were also interested in what types of imaginary AI technologies participants would develop.

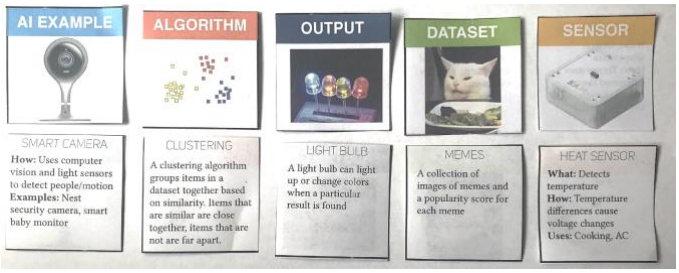

Figure 7: AI Example Card deck

5.3 Prototype Stations

In the second part of the workshop, participants engaged with the five low-fidelity prototypes of AI literacy exhibits—*Magic Mirror*, *LuminAI*, *Semantic Network*, *Neural Net*, and *Sensor Wall*. The goal of this stage of the workshop was to gather concrete feedback on specific ideas, in addition to familiarizing participants with the types of learning interventions and techniques used in museum spaces so that they could draw on these experiences in the final design creation stage of the workshop.

Each prototype station was set up with a design sketch of the future exhibit, a short write-up of the exhibit description, the prototype/activity materials, and a set of post-it notes and prompts for the participants to leave feedback. We asked the participants to respond to three feedback prompts for each prototype— "I like...," "I wish…," and "What if…?" This is a method that is often used to elicit prototype feedback within the interaction design community [78]. We purposefully brought low-fidelity, mostly paper-based exhibit prototypes in order to encourage participants to imagine future possibilities for design. Our intention was for each group to interact with at least three of the prototypes for ~15 minutes each. In reality, participants interacted with each station for 5-15 minutes because some stations—like *Semantic Network*—required longer/more complex interactions, whereas others—like *Sensor Wall*—were brief. Some groups also moved more quickly through the interactions, which is similar to how group interaction times and trajectories differ on the museum floor [80]. As a result, three groups (1, 3, 4) interacted with all five prototypes. Group 2 interacted with all prototypes except *LuminAI* and *Magic Mirror*; Group 5 interacted with all prototypes except *Neural Net*.

5.4 Design Activity

The final activity in the workshop was intended to be an opportunity for stakeholders to share their own novel ideas about designing AI literacy exhibits. One of the desired outcomes of the workshop was new ideas for design (in addition to feedback on our existing designs), and this was the focus of the culminating activity. Groups returned to their worktables after visiting the prototype station and were given several cards with prompts on them. Prompts were formulated as "Did you know…" questions with facts about AI related to the AI competencies [60] and "big ideas" [88], such as "Did you know that many AI devices are only able to do a specific task?" or "Did you know that some agents execute planned actions whereas others just react to objects in the environment?" These prompts were intended to scaffold the design experience and ensure that exhibit designs were focused on communicating factual, high-level concepts about AI. Participants were asked to design a museum exhibit to communicate the concept on the card. Drawing inspiration from the "bags of stuff" method that is often used in co-design [26], there were a variety of craft materials provided that participants could use to design their museum exhibits—including markers,

pencils, paper of different colors, scissors, glue, pipe cleaners, and other miscellaneous supplies like clay, googly eyes, styrofoam balls, etc. Participants were given ~25 minutes to complete the design task, with time left at the end of the workshop for participants to share their ideas with the group.

6 ANALYSIS

We took a mixed-methods approach to analyzing the data collected in the workshop due to the variety of different data types collected. In all aspects of our analysis, we avoid quantitative comparisons or statistics, since any quantitative results would not be statistically significant due to the small sample size of study participants. We focus instead on qualitative analysis, which is more well-suited to both our chosen methodology (co-design) and to identifying emerging patterns in a relatively new field of research [13].

We used a spreadsheet to track different components used and responses generated in the artifacts collected in the workshop (e.g. worksheets from the introductory activity, artifacts generated at the prototype stations), a method that is commonly used to make sense of "messy" data in alternative forms gathered in co-design workshops [78]. This analysis led us to identify several patterns in the artifacts, which we substantiate with qualitative descriptions of participant responses and creations. We also used the introductory worksheet packets to contextualize participant responses and ideas at the other activities based on their prior impressions and experiences. We recorded all participant feedback left at the prototype stations in a spreadsheet and identified common patterns. We transcribed all audio recordings and conducted an inductive thematic analysis of the transcripts, following the method outlined in [69]. We began by familiarizing ourselves with the data—a process which began during and after the workshop as we engaged with participants and took field notes. We continued this process by listening to all of the audio recordings, reading over our field notes, and jotting down points that stood out to us. We then reviewed our notes and identified common patterns.

After transcribing the audio recordings, we engaged in an initial round of line-by-line coding [1]. This involved coding each line of participant dialogue with an annotation describing the line (e.g. "Expressing interest in AI"; "Affirming participant A's suggestion"). Almost every line of dialogue was coded, with the only exceptions being a few lines with portions of inaudible dialogue or lines that researchers were unable to interpret the meaning of given the context. We additionally reviewed our field notes and the data from the collected design artifacts and added codes that augmented our findings from the audio recordings (avoiding redundancy). We took this approach in order to ensure a data-driven coding process. Since AI education for non-experts is a relatively new research field, we wanted to ensure that we were identifying new patterns that emerged from the data rather than relying exclusively on existing frameworks or our own potential biases for noticing patterns.

Next, we grouped the many (2,424) codes written in the line-by-line coding process into categories representing higher level themes (e.g. "Emotions about AI", "Mixed Feelings/Concerns about AI"). A number of these themes were novel; others emerged from the data but aligned with AI literacy competencies and design considerations [60]. We created a hierarchical structure of themes, with some more specific themes (e.g. "Privacy Concerns") belonging to higher level theme groups (e.g. "Mixed Feelings/Concerns about AI"). Some codes fit within multiple themes. The primary researcher kept a reflective journal throughout the process in order to keep track of how her impressions and ideas about the recordings changed over time [69]. The primary researcher was in charge of the codebook development, with incremental feedback from two other researchers. Conflicts were resolved via discussion until the researchers came to a mutual agreement. We describe our codebook and themes in more detail in the following section.

7 RESULTS

In this section, we report on the results from our analysis and discuss implications. The following three subsections correspond to the three stages of the workshop—introductory activity, prototype feedback, and design activity. A short description of each group that participated in the workshop is included in Table 2. All participant names used in this section are pseudonyms and ages are approximate.

7.1 Introductory Activity

7.1.1 Familiarity/Prior Experience with AI

All of the families demonstrated some level of *familiarity with AI* via the *AI example cards they discussed*, *the examples of AI they generated on their own*, or the *words they associated with AI*. Most of the adult participants and a couple of kids *referenced AI they had seen in popular media and science fiction* in addition to examples of AI they had previously interacted with. In some cases, experiences with popular media seemed to color participants' impressions of AI—for instance, Mark (G1) did not consider AI to be trustworthy since he had "seen *The Matrix*". Kimberly and Joseph (G5) made the distinction that they were not concerned about current AI technologies like "Alexa, or Google, or self-driving vehicles," but they were concerned about the types of AI they saw in movies ("like the one in *I, Robot*.")

Table 2: Descriptions of workshop participant groups

Group Number	Group Members	Notes
Group 1 (G1)	Mark (dad), Ella (daughter, age 6)	They were the only participants in the day 1 workshop. Both were actively engaged.
Group 2 (G2)	Maria (mom), Ricky (grandfather), Alex (son, age 10), Mimi (daughter, age 6)	Mimi was quiet through much of the workshop but actively participated in the design activity.
Group 3 (G3)	Susan (mom), John (son, age 16), Sam (son, age 16), Ellen (museum team member)	Group appeared to have some prior experience with AI/computing.
Group 4 (G4)	Joann (mom), Ralph (dad), Lauren (daughter, age 15-16), Katie (daughter, age 15-16)	Group appeared to have some prior experience with AI/computing.
Group 5 (G5)	Kimberly (mom), Joseph (dad), Liza (daughter, age 13-15), Anna (daughter, age 13-15), Kayla (daughter, age 7), Lucy (daughter, age 4)	Kayla and Lucy disengaged during some of the prototype station activities, and Joseph took them to play with some of the craft materials.

7.1.2 Mixed Feelings/Concerns about AI

All participant groups expressed *mixed feelings about AI*. Kids often expressed enthusiasm about AI, choosing words like *amazed*, *curious*, *interested*, and *hopeful*, although a few had reservations (Alex, G2; John and Sam, G3). Adults more often indicated *concern or hesitancy* (G3, G1, Maria, G2), though some also expressed hope (Kimberly, G5). All groups purposefully avoided circling *trustworthy* ("Do you trust AI? Not right yet." (Ralph, G4)). A few participants expressed that AI *confused or intimidated* them (G1, G2). Many participants *shared stories or anecdotes about past experiences they had with AI* as a way of expressing or explaining their feelings.

Several groups shared *concerns about AI*. G3 in particular discussed *privacy concerns*, *concerns about encoded bias*, and *AI taking people's jobs*. Several groups mentioned concerns related to *AI becoming self-aware or too powerful* (G1, G3, G5), although some were referring to fictional AI when they voiced this concern. Participants also discussed how *humans can take advantage of AI*, in scenarios ranging from news stories about hacking (G3) to anecdotes about using a Nest thermostat to mess with the temperature in the house behind a spouse's back (Maria,

G2). Concerns about AI sometimes revealed *conflicting feelings*—for example, Ellen (G3) liked using Alexa for convenience but was concerned about privacy.

Others expressed *skepticism about AI*. At one point during the workshop, Ricky (G2) asked incredulously, "There's people who actually go to school to study this?" Some participants reported purposefully avoided using AI technologies like social media (Mark, G1; Ellen, G3). During the introductory activity, G3 created an imaginary "nonsensical" AI that detected users' presence using a camera and looked at users' Facebook posts in order to decide whether or not to open a door. Sam (G3) commented on this imagined AI, stating that, "In a lot of ways, artificial intelligence, to me at least, seems like a gimmick…it's really trendy to put in machine learning, artificial intelligence on any…thing.'"

Table 3. Common themes from introductory activity discussions

<table>
<tr><td>1. Familiarity/prior experience with AI
1.1. Generating examples of AI
1.2. Discussing words associated with AI
1.3. Discussing AI example cards
1.4. Referencing popular media & science fiction</td><td>3. Dealing with AI errors
3.1. Noting that AI overlooks "edge cases"
3.2. Recognizing that minor alterations confuse AI
3.3. Suggesting human-in-the-loop systems to mitigate AI errors</td></tr>
<tr><td rowspan="2">2. Mixed Feelings/Concerns about AI
2.1. Confusion/intimidation towards AI/tech
2.2. Distrust/fear of AI
2.3. Enthusiasm and hope for AI
2.4. Interest in AI
2.5. Sharing stories of experiences with AI
2.6. Privacy
2.7. Hacking/humans taking advantage of AI
2.8. AI becoming self-aware or too powerful
2.9. AI taking jobs
2.10. Encoded bias
2.11. Skepticism about/avoidance of AI</td><td>4. Recognizing AI
4.1. Confusion about definition of AI
4.2. Associating AI with tech in general
4.3. Surprising examples of AI
4.4. Distinguishing between AI and humans</td></tr>
<tr><td>5. Imaginary Uses of AI
5.1. Household
5.2. School/work
5.3. Security/protection
5.4. Entertainment</td></tr>
</table>

7.1.3 Recognizing AI

Almost all of the groups (G1, G2, G3, G4) expressed some issues with *recognizing AI*—that is, confusion over distinguishing between artifacts that used AI and artifacts that did not. Several participants said they were *unsure what AI was* (Mark, G1; Ricky, G2) and others got into debates about whether certain technologies used AI (G3, G4). For instance, Ralph (G4) argued that GPS was not AI, it was "just a tracking system," but Lauren (G4) countered that "Google Maps now tells you if you have a traffic jam, if there's like a lot of people stopped at an intersection." Some of the groups (G2) suggested examples of AI that *conflated AI with technology in general*, like iPads or remote controls. Participants *found a few of the AI example cards surprising*, including the use of AI in video games (G1), online banking (G3, G5), and email spam detection (G3, G5). Kimberly (G5) was surprised that AI was used to recognize checks because she assumed "a person [was]...taking care of it." Some participants *drew distinctions between AI and humans* ("[AI is] something that isn't a person but knows how to think" (Maria, G2); "We're intelligent, like we're thinking people because we have brains in our heads. We're talking about artificial meaning, it's not natural. It's made. We made it." (Kimberly, G5)).

7.1.4

7.1.4 Imaginary Uses of AI

When prompted to use the provided input, algorithm, and output cards to create an imaginary AI, groups explored a variety of different *imaginary uses of AI* including AI that could be *used for entertainment* (Ella, G1) or *protection/security* (G5). Several groups imagined AI devices that could help them with *household tasks* (G2, G5), such as taking out the trash or making lunch (G5). Other groups suggested having AI help them with *school/work related tasks* like taking a test for you (G2) or doing your job (G3). This is consistent with a previous workshop study where the authors found that when asked to imagine future AI, kids often created AI that could help them with tasks like household chores or homework [25]. However, some groups in our study were also skeptical about having AI do tasks for them—Ellen (G3) in particular was concerned that if AI did her job for her, she would be replaced.

7.1.5 Dealing with AI errors

Several participant groups considered different errors that AI might make and how to deal with them when thinking of their imaginary AI. John and Sam (G3) recalled *'edge cases'* that cause AI to fail, like their phone's inability to tell their faces apart since they are twins. G3 and G5 both recognized that *minor alterations can often confuse AI*. For instance, G5 designed an AI security system that operated using facial recognition and considered issues like how the AI would know not to call the police on the mailman dropping off mail, their grandparents coming for a visit, or people attending a party they were throwing. Kimberly suggested a *human-in-the-loop* phone alert system where "it would just alert me that a stranger is walking up...and I would be like, it's the mailman."

7.2 Prototype Interactions

This section discusses participant interactions with the exhibit prototypes. We present themes related to social interaction dynamics (i.e. whether/how group members collaborated, including commentary on the division of labor amongst group members [19]), social learning talk [4,76] related to the exhibit, and participant feedback on the exhibit design. We see both the discussion of social interaction dynamics and social learning talk as being relevant to understanding how learners collaborated at the exhibits (i.e. engaging in verbal dialogue, shared interactions, and/or collectively working towards a shared goal). We identify participant comments where possible, but feedback post-it notes were anonymous so some quotes are presented without a pseudonym. Key themes are presented in tables and discussed in more detail in the text.

7.2.1 Social Interaction Dynamics

Different social interaction dynamics emerged at each of the exhibit prototypes. Learners tended to explore the *Sensor Wall* and *Neural Net* activities together as a single group. At *Magic Mirror* and *Semantic Network*, some groups collaborated on a single artifact (cyborg/network), while other groups split up and made multiple artifacts. There was a clear division of labor between adults/older children and younger children at both *Magic Mirror* and *Semantic Network*. For instance, at *Magic Mirror*, older participants typically read the cards aloud while the younger kids participated in the conversation and glued cards onto the outlines. At *LuminAI,* group members generally took turns dancing—or nominated one person to try out different moves while the rest of the group looked on. Similarly, when interacting with *MoViz*, group members usually stood around the screen and commented on the visualization while one person navigated around the UI.

Table 4: Summary of social interaction dynamics at each exhibit prototype

Exhibit Name	Group Dynamics	Division of Labor
Magic Mirror	Single or multiple groups	Older group members read cards, younger kids discussed and glued cards
Sensor Wall	Single group	None
Neural Net	Single group	None, although some younger kids disengaged
Semantic Network	Single or multiple groups	Older group members planned network, younger kids identified cards and glued
LuminAI	Learners took turns dancing with *LuminAI* or interacting with *MoViz* while others observed	One person was the active user, others took on role of commentator/observer/director

7.2.2 Social Learning Talk

Participants engaged in a variety of social learning talk while interacting with the exhibits. Roberts et al. describe four different types of learning talk that commonly occur at museum exhibits—*manage* statements are "related to the establishment of joint attention, negotiation of action, or scaffolding exhibit use"; *instantiate* statements "indicate…when a user says aloud a piece of information, providing opportunities for other visitors to internalize that information"; *evaluation* statements "make a judgment or assessment about a piece of information by assigning some kind of value, whether qualitative or quantitative"; *integration* is "the act of pulling together multiple pieces of information presented in an exhibit"; and finally *generate* statements "combine information from the exhibit with visitors' own prior knowledge and experiences" [76]. We loosely structure our discussion of learning talk at the exhibit prototypes according to this framework (keeping in mind that a single statement can often contain multiple different types of learning talk).

Management talk occurred at all of the exhibits. We particularly saw a lot of management at *LuminAI*, where participants took turns dancing with *LuminAI* while others directed them to try out different moves. When interacting with *MoViz*, the cluster visualization tool, participants directed the participant in charge of the controls to navigate around the interface ("Click on the blues and make it appear." (Ella, G1); "I want to see if you can move it. What if you click where there's a bunch of them, what does it do?" (Joann, G4)).

Instantiation played a particularly important role at *Magic Mirror*. Parents and older kids read sensor/actuator cards aloud, allowing the whole group (including younger, pre-literate children) to internalize information about the sensor or actuator. Instantiation also played a role at *LuminAI*, where some participants read aloud the text on the screen. The text prompted users to teach *LuminAI* and explained what kind of moves the dancer was performing (e.g. "Cool move! I'm going to do something similar.") Reading this aloud may have helped learners to internalize that the dancer was learning from them and consider some of its decision-making modes.

Evaluation statements were common at several of the exhibits. Participants (G1, G3, G4) qualitatively evaluated particular gestures in *MoViz* ("Well look at this one–so I have it on this little blue dot and all of a sudden the arms start flailing!" (Susan, G3)) and observed similarities between gestures that were in the same cluster ("Blue's really complex. See like look how its shifting the legs and also kinda doing arm motions and then red is just like, it's just kinda arms." (John, G3)). The most common interaction with *LuminAI* was physically trying out different moves and vocalizing/commenting on the moves *LuminAI* was performing in response on the screen. Some participants tried to find bugs or test edge cases by interacting with *LuminAI*. For example, Susan (G3) investigated whether the system would break if two people were in the Kinect sensor's field of view at once by testing it out with her son Sam. Several users also commented on the limitations of the Microsoft Kinect motion sensor.

Similar interactions also occurred at *Sensor Wall*, where participants tested out a variety of different sounds to see how the exhibit would respond, including clapping, snapping, whistling, beatboxing, singing, and speaking at different volumes, pitches, and durations. Participants also made observations as to what happened when they made the sounds (i.e. noticing the screen blink or the LED light on the Arduino light up). Some participants pointed out some limitations of the sound sensor ("It's not that sensitive" (Sam, G3), "I feel like it picks up the high pitch noises better" (Mark, G1)).

At *Magic Mirror,* participants (G1, G4, G5) engaged in evaluation when they qualitatively compared and contrasted human body parts to sensors/actuators and recognized that several human body parts could not be one-to-one replaced with machine components—for instance, Joseph (G5) noted that the mouth "does more than just...taste. 'Cause it talks too or produces sound," and noted that you would have to use several sensors and actuators to fulfill that purpose. Participants also engaged in qualitative assessment of different creatures at *Neural Net*, discussing the features of the animals on the example cards (e.g. "it [an iguana] doesn't have feathers" (Ella, G1)). Younger participants in particular looked through the cards and guessed which ones they thought would be classified as birds (Ella, G1; Alex, G2).

Integration, or making connections between exhibit components, was particularly common at *Semantic Network*. Participants discussed where to place concepts in their network and what relationships to connect concepts with (e.g. "Does your kitty have a nose?" (Maria, G2)). In doing so, they found commonalities between the different concepts ("Ooh dogs have whiskers right? So do cats." (Mark, G1)). Some participants noted that it was difficult to express certain nuanced concepts in the network (e.g. "Cats dislike birds. Right? Probably? No, they eat them." (Mark, G1)).

Participants integrated information at *Neural Net* when they discussed how they thought the weights would affect the classification of the examples (e.g. "But if you focus, if you keep showing it pictures when you're training it of beaks, and then you show it a platypus, it's going to think it's a bird." (Joann, G4)). G3 participants were interested in how multiple weighted features would interact and wanted to test out different combinations ("imagine, like, you have maybe a heavy weight on like skeletal structure, maybe a beak, but like...not a lot of weight on eggs." (Susan, G3)).

We saw the most *generate* statements, or statements connecting with learners' prior knowledge/experiences, at *Magic Mirror* and *Semantic Network.* At *Magic Mirror,* most groups began the process of adding parts to their cyborgs by thinking about sensors and actuators that were analogous to human body parts ("This detects sound. Now what detects sound on your body?" (Joseph, G5)). This tied in with our design goal of having participants draw on their "body knowledge" as they made sense of different sensors and actuators. Groups also engaged in learning talk surrounding several prompt questions that were on the table, including considering which of their own body parts they would or would not replace with sensors. Kimberly (G5) did not want to move the location of her body parts around, referencing religion when she said that "I don't want my eyes down here. I'd want them where they are. I want ears on the sides of your head so you can hear the most, or at least sound sensors can hear the

most...sound," because "I think God did that on purpose." Ella (G1) was excited by the prospect of placing body parts, and was particularly interested in the interface between the human body and machine body parts, like what it would be like to lose a robot tooth or how you would trim a robot toenail. Several participants wanted to use sensors to either fix their own body parts or enhance their natural capabilities. For example, Anna (G5) suggested replacing her ears because "when you grow older you lose your sense of hearing." Participants also considered things that could go wrong when replacing body parts with sensors. For instance, Mark (G1) wondered about what would happen if he replaced his eyes with cameras and they malfunctioned.

Participants also drew on prior knowledge frequently at *Semantic Network*. Several participants chose a "theme" for their network before getting started—many of these were based on personal interests. For example, G1 decided to start their network with a dog card because they had dogs. G5 themed their network around musical instruments they played. Participants also related the semantic network to other similar structures they had seen before—like a "chart" (Ella, G1) a "web" (Mark, G1), or brainstorming/making a collage (Ricky, G2). Several participants also made the connection that building the semantic network was like teaching or programming a computer (Maria, G2, Mark, G1).

Table 5: Summary of learning talk at each exhibit prototype. A blank cell does not indicate that no learning talk of that nature occurred at the prototype, just that it was not significant enough to constitute a key theme in our analysis.

Exhibit Name	**Manage**	**Instantiate**	**Evaluate**	**Integrate**	**Generate**
Magic Mirror	Discussing parts to add to cyborg	Reading cards aloud	Comparing human body parts with sensors / actuators	-	Considering parts to replace; drawing on body knowledge
Sensor Wall	-	Noticing sound or visual	Trying out different noises and making observations about system reaction to sounds; pointing out limitations of sensor	-	-
Neural Net	-	-	Discussing features of animals; Guessing what the animals would be classified as	Considering how weights would affect classification	-
Semantic Network	-	Reading cards aloud		Discussing relationships; noticing commonalities	Choosing personal theme for network; comparing network to other structures
LuminAI	Directing others to dance or navigate *MoViz* UI	Reading text on screen aloud	Trying moves and observing response; noticing gestures and similarities; searching for bug; discussing Kinect	-	-

7.2.3 Design Feedback

Participants had positive feedback for exhibits that allowed for open-ended creative interactions—particularly *Magic Mirror* and *Semantic Network*. Learners liked how they could customize their creations ("I liked that you got

to show…how you would want your body" (Ella, G1); ("[*Semantic Network* was] fun because I got to do whatever I wanted." (Alex, G2)). Participants also liked how the prototypes allowed them to be creative and draw on their prior knowledge ("I liked seeing how the sensors are compared to our body parts" (Liza, G5); ""I like the different options and connections made between very different things [at *Semantic Network*]. It's like a puzzle.")

Table 6: Summary of key participant design feedback themes for each prototype

Exhibit Name	Design Feedback
Magic Mirror	Enjoyed open-ended hands-on nature; liked drawing on "body knowledge"; wanted to engage with real sensors or connect with *Sensor Wall*
Sensor Wall	Enjoyed responsiveness and interacting with a real sensor; many thought interaction was too simple and wanted to expand it; suggestions to connect with *Magic Mirror*
Neural Net	Suggestions to integrate technology, make it more accessible for young learners, add connections to the "real world" and *Semantic Network*
Semantic Network	Enjoyed open-ended hands-on nature; suggestions to integrate technology and connect with *Neural Net*
LuminAI	Enjoyed dancing; older participants enjoyed exploring *MoViz*; suggestions to integrate *MoViz* and *LuminAI*

Exhibits that provided immediate, "visceral" [83] responses also received positive feedback. Several groups commented on how they liked the "immediate feedback" from *Sensor Wall* and appreciated how it responded to noises (Ella & Mark, G1; Maria, G2). Similarly, participants enjoyed dancing and "moving around" with *LuminAI* (Sam and John, G3; Mark, G1). Suggestions for expanding exhibits also focused on improving responsiveness or adding more interactive features. *Sensor Wall* in particular elicited a lot of suggestions for expansion—while some groups appreciated the simplicity of the experience, others felt it was too simple and all groups suggested making the visual response more interesting or adding additional sound recognition features. Several participants suggested expanding the *Neural Net* activity to make it more responsive, using lights and visuals to indicate the weights on certain features (G3) or highlighting relevant features in the bird images (Ella, G1). Learners also had suggestions for making *Semantic Network* more responsive—e.g. "I wish there was an animated AI that built what it thought we meant in real time as we placed the pieces down." Several groups that interacted with *Magic Mirror* wanted to try out the sensors that they incorporated in their cyborg to see how they actually worked (G1, G3, G5). A common theme across these comments was the desire to interact with something "real"—a real sensor, a "real neural network," or a real AI using the semantic network.

Most groups felt that it was important for exhibits to be accessible to learners of all ages, and some were frustrated when the exhibit was hard for their youngest kids to engage with. Learners liked that *Magic Mirror*, *Semantic Network*, and *Sensor Wall* were engaging for all group members ("I like that you can do it with your children"). On the other hand, numerous participants (G2, G5) commented that *Neural Net* was a little difficult to understand, particularly for younger kids, and suggested providing clearer instructions, allowing kids to "to practice with trial and error to see what the best input would be," or (as mentioned) incorporating interactive components like "lights and sounds" to "make it more fun."

Participants had numerous suggestions for how to make connections—between exhibits, to other disciplines, and to the "real world." Several groups suggested connecting multiple exhibits into one larger experience—for instance, connecting *Magic Mirror* with *Sensor Wall* (Sam, G3) to enable participants to try out the sensors in their cyborg, or connecting *Semantic Network* with *Neural Net* (Ellen, G3) to communicate the difference between top-

down programmed knowledge and bottom-up learned knowledge. Numerous groups wanted to connect *MoViz* with *LuminAI* in a more meaningful way—for instance, coloring the dancer's gesture according to which *MoViz* cluster it belonged to (Mark, G1). Several participants also had ideas for how to *connect exhibits to other disciplines.* Ellen (G3) thought that *Neural Net* could be used to teach young kids about categorization, and Susan (G3) suggested using the weights in *Neural Net* to teach about math percentages. Several participants wanted to strengthen the connection between *Neural Net* and "real life." Participants suggested exploring alternative domains for the neural network like facial recognition, racial discrimination, or climate change (Susan & Ellen, G3).

7.3 Design Activity

This section describes the artifacts generated by participants when they were prompted to create an exhibit to teach visitors about AI. These artifacts can provide insight into visitors' priorities and interests in AI-related exhibits, which can then inspire future design research. In this section, we describe the exhibits that participant groups designed. We comment on common design features in Discussion (section 8). We include a few participant-drawn sketches where they add to an understanding of the textual description of the designed exhibit.

7.3.1 Group 1: AI Paintings & Riddles

Mark (dad) and Ella (age 6) selected the prompt card "Did you know that some AI can do creative things, like write poems or generate paintings?" Ella wanted to create an exhibit that showed that AI could write poems and Mark wanted to create an exhibit that showed AI drawing a picture, so they decided to each create a separate design. Ella took inspiration from one of her favorite exhibits at the museum, which has an auditorium with a spherical projection screen in the middle. She suggested having the screen display a riddle ("question poem") with multiple choice answers for the audience to choose between. The visitors sitting in the auditorium would have buttons or clickers that they could use to choose an answer to the riddle that was generated by the AI on the screen. Mark situated his exhibit in a hallway. Visitors walking down the hallway could select between different features of a painting—for example, choosing between a landscape vs. a portrait, or selecting color vs. black and white. Based on these inputs, the AI would generate a unique painting that it would project on the wall in the hallway. Visitors could view a "time-lapse" video of the AI generating the painting. Visitors would be able to email themselves a copy of the picture they created.

7.3.2 Group 2: Drone Showcase

Group 2 picked the prompt, "Did you know that AI devices use sensors to detect things in the world?" Alex (age 10) suggested creating a "drone showcase" exhibit where you could fly a drone around the museum (he was really enthusiastic about drones throughout the workshop). Mimi (age 6) came up with the idea to move drones around in a room to solve a puzzle. This reminded Alex of "mini games, where you have to...get an egg and you have to fly it

to the other side of the room without it cracking." Ricky suggested the idea of using the drone to walk the dog. The group ended up sketching an exhibit where a visitor could fly a drone to bring food and water to a dog.

7.3.3 Group 3: Competing Robots

Group 3 chose the prompt, "Did you know that some robots move by pre-planning routes, and others simply react to objects in their path?" The group came up with the idea of an exhibit in which two different robots would navigate a maze. One robot would navigate the maze by sensing obstacles and the other would operate according to a visitor-written program. Visitors could use a block-based programming language on an iPad kiosk to program the pre-planned robot. They could then compare the performance of the two robots and discuss the pros and cons of each approach.

Group 4: Responsive Box

Group 4 selected the design prompt, "Did you know that some AI systems are often only good at one very specific task—like recognizing a face or turning on a light when motion is detected?" Lauren suggested creating "a little dude that runs away when you put your hand near it." She suggested using a motion sensor to detect the presence of your hand. Ralph added to this idea, suggesting using speed to modulate the response of the AI. They ended up creating a moving box that sensed hand proximity and speed. If a hand approached it quickly, the box would move away. If it approached slowly, the box would move towards you, and if you touched it, a heart would show up on a touchscreen. If you moved your hand towards the box at medium speed ("just right!") it would stay still. If you touched it, you would see two hearts.

7.3.4 Group 5: Family-Friendly Smell Detector

Group 5 picked the prompt, "Did you know that humans play an important role in fine-tuning AI systems?" They designed a "family friendly" exhibit that "little kids would think was fun too." Visitors would step into a booth where an AI would learn how "humans think different chemicals and pheromones smell." The exhibit would:

> "shoot a little scent out and then it would, for the sake of like smaller or younger children, have maybe like a touchscreen with pictures and they could choose between certain things like what it smelled like. So, like tap a picture. And then like a secondary question would come up that said, 'is this a pleasant smell or is this an unpleasant smell?' So, it can help the smelling sensor kind of figure out what we think and how we feel about certain scents" (Kimberly, G5).

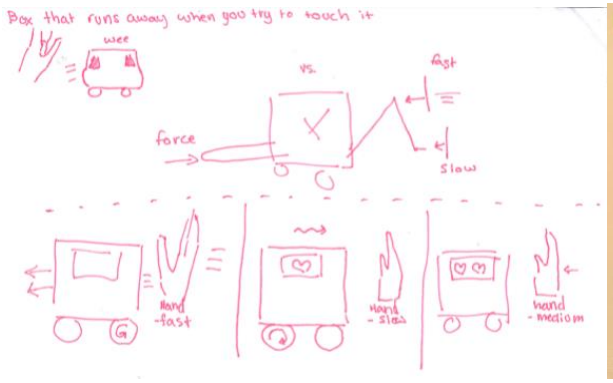

Figure 8: Diagram and physical prototype of G4's responsive AI box exhibit (left); G5's family-friendly smell detector exhibit (right)

8 DISCUSSION

Themes that we found in our analysis reveal some insights related to designing exhibits to promote AI literacy. In this section, we reflect on how our core design considerations—*collaboration, embodied interaction,* and *creativity*—worked in practice at each exhibit prototype. We also present design recommendations for future AI literacy exhibits based on the themes we noticed in the workshop discussions. Several of these recommendations echo design considerations for AI literacy learning interventions [60] and are in conversation with prior work on collaborative learning (e.g. [35,38,50]).

8.1 Core Design Considerations

8.1.1 Collaboration

All exhibit prototypes supported some degree of collaboration between group members, although the types of collaboration that emerged differed by exhibit (e.g. *Semantic Network* and *Magic Mirror* facilitated more collective action as participants worked towards the shared goal of creating a network or cyborg; *Neural Net* facilitated more shared dialogue). Features of participant generated exhibit designs also suggested that participants viewed collaboration as being important. G1 and G3 specifically designed their exhibits to support social interaction—Ella (G1) designed an exhibit that engaged audience members in a collaborative Q&A session, and G3 designed a game-like exhibit where multiple visitors could work as a team to program a robot to navigate a maze.

Participant feedback on exhibit prototypes indicated that iterative changes could be made to our plans for exhibit design to improve support for collaboration. For example, enabling additional user shadows in *LuminAI* could allow for more shared dance experiences and collective action, and keeping *MoViz* on a screen or projection rather than putting it in VR as we originally intended could help to foster some of the shared dialogue about gesture clusterings we saw in the workshop at the final exhibit. Similarly, participants enjoyed collaborating on their cyborgs in *Magic Mirror* (collective action), an activity we likely want to consider incorporating into the final exhibit (rather than just projecting sensors onto the mirror). The prompt questions included with *Magic Mirror* also spurred a lot of social learning dialogue, and we may want to include similar questions for the other exhibits in the future. Providing multiple interaction stations for activities like *Semantic Network* could allow multiple subgroups to create artifacts simultaneously [46,49]. Additional steps can also be taken to better engage young learners (see 8.2.6 Support for young learners).

The findings from this study also indicate more broadly the importance of incorporating opportunities for collaboration in AI learning experiences. AI is a relatively new topic to many people, it is complex, and there are many open questions about its societal and ethical implications. Collaborative social dialogue can aid learners in constructively debating and making sense of AI and the issues surrounding it [23]. We saw this occur in the Introductory Activity, when personal narratives and multi-generational dialogue emerged as family members shared their different experiences and feelings towards AI technologies. This type of shared dialogue is important to foster if we want learners to engage in critical discussion about AI and its role in their personal lives.

In addition, AI (and its association with computer science) can be characterized as an individualistic pursuit [6] or as cold/mechanical [71,85]. Engaging in social learning activities about AI can potentially aid in making the field more engaging and interesting for a more diverse audience. Learners seemed to engage in the most learning

dialogue at the activities where they were working together towards a shared goal—*Magic Mirror* and *Semantic Network*—rather than just engaging in dialogue surrounding a shared activity.

8.1.2 Embodied Interaction

All of the exhibits also engaged participants in some form of embodied interaction [51]. Participant feedback illustrated that participants liked the "hands-on" nature of the exhibits or enjoyed activities like "finding pictures," "moving pieces," or "pressing buttons." All of the participants' design sketches involved interactive activities including using buttons to provide visitor input (G1), programming and physically playing with robotic devices (G3, G4), teaching an AI about smells using a touchscreen kiosk (G5), and flying a drone (G2). This suggests that participants are excited by tangible experiences, which may stem from their existing expectations about museums and/or the abstract nature of AI being made concrete. Prior work suggests that adding in more interaction stations [46], tools for exploring sensors' capabilities [49], and building in opportunities for reflection [25,38] could all potentially lead to increased learning talk.

8.1.3 Creativity

Only three of the prototypes supported open-ended creative experiences. Participants were able to engage creatively with *Semantic Network* and *Magic Mirror*, building networks and cyborgs related to their personal interests. Participants were also able to engage creatively with *LuminAI*, exploring how to co-create dance moves with the AI, although we did not observe as much co-creation during the workshop as we have during past installations of *LuminAI* in public spaces [58]. This may have been partially due to situational issues, like participants being embarrassed to dance in front of people without music in the small workshop room or *LuminAI* being slow to respond due to laptop issues on the second day of the workshop. Overall, our results indicate that **the exhibit prototypes that supported all three design considerations during the workshop—*Semantic Network* and *Magic Mirror*—fostered the most learning talk amongst participants of all ages.** This suggests that the design considerations driving our exhibit designs—*embodied interaction*, *collaboration*, and *creativity*—were effective at providing engaging experiences that promote learning talk, **especially when used in combination with each other**.

Participant design sketches and prototype feedback also supported the idea that creativity was an important design consideration. Several participants suggested adding creative components to *Sensor Wall*, like making the sensor respond to input with visual art or music. Participants also asked for the ability to co-create with each other (in addition to the AI dance partner) at the *LuminAI* exhibit. G1's design sketches focused on AI that could creatively generate riddles and paintings, and G3 incorporated a creative visitor programming activity into their maze exhibit.

8.2 Additional Design Considerations

8.2.1 Addressing issues with recognizing AI

A common theme in the introductory activity was confusion about what AI is and what technologies use AI. Even after completing the introductory activity, there was confusion about the differences between technology in general, algorithms, and AI. Defining AI is a foundational topic, but one that even experts debate [60]. Considering the importance of this issue for being able to engage with AI in everyday life, designers of AI literacy learning interventions may want to consider prioritizing engaging learners in discussions about what AI is and how to recognize it.

8.2.2 Building on learners' prior knowledge and interests

Prior work suggests that learners will likely bring preconceptions about AI to AI-related learning experiences, many of which come from popular media like movies, television shows, or news media [60]. This contention is supported by our findings—most of the adults in the workshop referenced AI in popular media. Although a few children referenced popular media, they tended to draw more on their own prior experiences interacting with AI technologies. Some of the participant preconceptions we discuss in Results (section 7) could be leveraged in future design interventions to create learning experiences that build more effectively on participants' prior knowledge. We suggest that exhibits that connect to commonly used AI technologies may help young learners in particular to draw on their prior experiences when making sense of AI in addition to helping them to better understand the technologies they use in their daily lives. This has not been the focus of past museum exhibits that have exhibited unusual robotic devices that are not widely used outside of educational settings or research labs. Incorporating technologies that kids are enthusiastic about—like video games (Lauren, G4; Alex, G2), music (Ella, G1), or drones (Alex, G2)—may also help to encourage engagement. If incorporating references to popular media like TV or movies, designers may want to consider generational differences—older movies or television shows may resonate with adults but may not be as relatable to kids.

Future research investigating preconceptions about AI in a larger, more representative audience is necessary and could lead to important insights for designers and educators. In particular, investigating how AI literacy learning interventions affect preconceptions and feelings about AI (including interest) is a promising area for future research. Designers may also want to consider engaging in further co-design work to investigate how to incorporate children's existing "funds of knowledge" [35] about AI in novel museum exhibits—for example, Iversen and Dindler conducted a workshop in which they encouraged children to think about qualities of games and social sites that they use that make those platforms fun and engaging, then worked with the children to apply these qualities to a mock-exhibition space [20]. Similar work could explore how to incorporate aspects of AI that children find interesting in exhibit design.

8.2.3 Incorporating emotion and narrative

Designers may also want to consider tapping into some of the mixed emotions that participants have surrounding AI. Exhibits that spur discussion and/or constructive debate [23] about AI-related ethical issues like privacy, bias, or the future of work could help learners to collaboratively make sense of issues they may have already grappled with in their personal experiences. Sensitive or controversial topics like AI's relationship to religion (G5's commentary during *Magic Mirror*) and violence (Alex (G2) mentioned drones dropping bombs and sensors in landmines during the design activity) came up without prompting during the workshop. Carefully designed learning experiences could foster more extended discussion about such topics, potentially providing opportunities for crossover learning about other disciplines like ethics or philosophy. Encouraging learners to discuss their feelings about AI technologies is also important as emotions and feelings are a key aspect of learning, particularly in informal spaces [10,11]. Narrative and storytelling can play an important role in these types of discussions [11], as we saw in the workshop when many participants shared anecdotes about both positive and negative prior experiences with AI. Designers may want to consider explicitly encouraging learners to share stories about their prior experiences. Prior research in the CSCW community on collective storytelling and generating shared imaginaries of future technologies (e.g. [64]) could inform future work in this design space.

8.2.4 Real world relevance

Participants expressed interest in how AI technologies could be related back to "real-world" issues, and were generally enthusiastic about opportunities to interact with "real" AI technologies (see 8.2.7 Integrating technology). Exhibits that draw on personal interests and easy-to-grasp concepts like family, animals, etc. can serve as an easy entry point (particularly for younger kids), but this emergent theme indicates that designers should not shy away from addressing real/current issues with AI. Discussions of "real world" issues can also tie in with discussions of ethics and draw on learners' existing "funds of knowledge" [35], as mentioned in 8.2.2 Building on learners' prior knowledge and interests.

8.2.5 Making connections

Several participants suggested drawing connections between the exhibits that they engaged with in the workshop. One participant (Mark, G1) suggested making connections between the museum learning experience and interactions at-home by incorporating the ability to "take home" an AI painting in his exhibit design. This common theme of wanting to build connections suggests that designers of AI learning experiences should not consider their learning interventions in isolation but rather make an effort to draw connections between exhibits (and potentially other disciplines) so that learners can gain a more holistic picture of AI. Designers may also want to consider helping visitors to connect their museum learning experience with other learning environments like the classroom and at-home learning experiences. We are starting to explore this by developing lesson plans and classroom/at-home activities related to the exhibit prototypes. Ito et al.'s approach to designing connected learning experiences [50] can help guide the design of experiences that extend the museum experience into classrooms or at-home learning spaces.

8.2.6 Support for young learners

Most groups expressed concern about the ability of exhibits to support young learners. Several groups discussed how exhibits could be enhanced to support young children in addition to adults and older children, suggesting more visuals/pictures instead of text (G2, G3, G5), content related to early childhood development (G3), incorporating humor (G2), and multiple levels of entry supporting family-friendly interactions (G5). G5's exhibit design was specifically geared towards supporting engaging interactions for younger kids. It is worth noting that children at the low end of our target age range (particularly Ella (G1, age 6)) were able to engage in learning talk surrounding the exhibit prototypes and workshop activities. Even younger kids seemed to enjoy visceral AI-related experiences like playing with responsive sensors or scaffolded experiences where it was easy for them to work together with their older family members. Designers should consider how to create exhibits that can support learning about AI amongst learners of all ages.

8.2.7 Integrating technology

Participants engaged in learning dialogue about AI at paper-based exhibit prototypes (*Magic Mirror*, *Semantic Network*, *Neural Net*) despite not actually interacting with AI technologies. However, participants were also enthusiastic about future plans to incorporate technology in these exhibits. In addition, all of the participants' design sketches involved interaction with actual AI devices like robots, drones, sensors, and machine learning. Sensors and actuators in particular seemed to elicit a lot of participant interest/dialogue throughout the workshop. The success of the paper-based prototypes at facilitating learning dialogue is promising for the development of "unplugged" (i.e.

no computer needed) [55,68] AI learning experiences for low-cost learning interventions. However, participant feedback indicated that—particularly in museum spaces—allowing for interactions with novel AI technology would not only be engaging and potentially inspire awe [11,52] but could also facilitate opportunities for active reflection on learner-generated artifacts like cyborg designs or semantic networks. This is in line with prior AI education research that suggests that receiving immediate feedback is important in AI learning experiences [25,98], as well as research on exhibit design more broadly that emphasizes the importance of incorporating opportunities for reflection and metacognition [38].

9 CONCLUSION

In this paper, we investigate how to design museum installations to improve public understanding of and interest in AI. We present the design of five exhibit prototypes—*Magic Mirror, Sensor Wall, Neural Net, Semantic Network, and LuminAI*—that aim to communicate a variety of different AI literacy competencies. We then present the results from a co-design study that we conducted, presenting themes related to participant preconceptions about AI, learning talk relating to AI literacy competencies at exhibit prototypes, and design features that participants deemed important in their feedback on existing prototypes as well as in their own unique exhibit designs. Our findings suggest directions for future research in understanding the effects of preconceptions on learning and interest development in AI, new topics for AI-related exhibits that may address participants' interests, concerns, and spark enthusiasm, and design considerations for developing AI literacy exhibits in informal spaces. The five exhibit designs that we present in this paper can also serve as inspiration for other researchers looking to develop informal learning experiences surrounding AI for family groups or K-12 students. This paper contributes to the CSCW community by 1) presenting five exhibit prototypes that incorporate collaboration as a key design goal; 2) exploring the types of social learning dialogue that emerge from interactions with AI literacy exhibits and the role that collaboration plays in learning abnout AI; and 3) utilizing co-design practices to explore family groups' interests and priorities with regards to AI literacy learning experiences.

ACKNOWLEDGMENTS

We would like to thank Jonathan Moon, Lucas Liu, and Cassandra Naiomi for their contributions to the workshop materials and prototypes; Dr. Betsy DiSalvo for her advice and feedback; Sara Milkes, Erin Truesdell, and Dr. Hyunjoo Oh and her Designing for Curiosity course for helping to pilot test workshop activities; Patrick Fiorilli and Sucheta Ghoshal for paper feedback; and our collaborators at the Museum of Science and Industry, Chicago—especially Dr. Aaron Price and Natalie Harris—for their ongoing support of this research. This work was funded by the National Science Foundation (DRL #1612644).

REFERENCES

[1] Dor Abrahamson and Arthur Bakker. 2016. Making sense of movement in embodied design for mathematics learning. *Cognitive research: principles and implications* 1, 1: 33.

[2] Edith Ackermann. 2004. Constructing knowledge and transforming the world. *A learning zone of one's own: Sharing representations and flow in collaborative learning environments* 1: 15–37.

[3] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives* 31, 2: 211–36.

[4] Sue Allen. 2004. Designs for learning: Studying science museum exhibits that do more than entertain. *Science Education* 88, 1: S17.

[5] Sue Allen, Patricia B Campbell, Lynn D Dierking, Barbara N Flagg, Alan J Friedman, Cecilia Garibay, and David A Ucko. 2008. 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*.

[6] Morgan G Ames. 2018. Hackers, Computers, and Cooperation: A Critical History of Logo and Constructionist Learning. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW: 18.

[7] Roya Jafari Amineh and Hanieh Davatgari Asl. 2015. Review of constructivism and social constructivism. *Journal of Social Sciences, Literature and Languages* 1, 1: 9–16.

[8] Alissa N Antle, Greg Corness, and Allen Bevans. 2013. Balancing justice: Comparing whole body and controller-based interaction for an abstract domain. *International Journal of Arts and Technology* 6, 4: 388–409.

[9] Ars Electronica Futurelab. 2019. *Understanding AI*. Retrieved from https://ars.electronica.art/aeblog/en/2019/08/06/understanding-ai-futurelab-installations/

[10] Constance Barlow and Barry L Hall. 2007. What about feelings?: A study of emotion and tension in social work field education. *Social Work Education* 26, 4: 399–413.

[11] Leslie Bedford. 2014. *The Art of Museum Exhibitions: How story and imagination create aesthetic experiences*. Left Coast Press.

[12] Margaret A Boden. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge, New York, NY, USA.

[13] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2: 77–101.

[14] Leah Buechley, Mike Eisenberg, Jaime Catchen, and Ali Crockett. 2008. The LilyPad Arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 423–432. https://doi.org/10.1145/1357054.1357123

[15] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 77–91.

[16] Timothy Charoenying. 2013. Graph hopping: learning through physical interaction quantification. In *Proceedings of the 12th International Conference on Interaction Design and Children*, 495–498.

[17] Nikolaus Correll, Chris Wailes, and Scott Slaby. 2014. A One-hour Curriculum to Engage Middle School Students in Robotics and Computer Science using Cubelets. In *Distributed Autonomous Robotic Systems*. Springer, 165–176.

[18] Kate Crawford and Trevor Paglen. 2019. *Training Humans*. Retrieved from http://www.fondazioneprada.org/project/training-humans/?lang=en

[19] Joshua A Danish, Noel Enyedy, Asmalina Saleh, and Megan Humburg. 2020. Learning in embodied activity framework: a sociocultural framework for embodied cognition. *Int. J. Comput. Support. Collab. Learn.* 15, 1: 49–87.

[20] Christian Dindler, Ole Sejer Iversen, Rachel Smith, and Rune Veerasawmy. 2010. Participatory design at the museum: inquiring into children's everyday engagement in cultural heritage. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction*, 72–79.

[21] Daniella DiPaola, Blakeley H Payne, and Cynthia Breazeal. 2020. Decoding design agendas: an ethical design activity for middle school students. In *Proceedings of the Interaction Design and Children Conference*, 1–10.

[22] Zachary Dodds, Lloyd Greenwald, Ayanna Howard, Sheila Tejada, and Jerry Weinberg. 2006. Components, curriculum, and community: Robots and robotics in undergraduate AI education. *AI magazine* 27, 1: 11–11.

[23] Willem Doise, Gabriel Mugny, and Juan-Antonio Pérez. 1998. The social construction of knowledge: Social marking and socio-cognitive conflict. *The psychology of the social* 77.

[24] Paul Dourish. 2004. *Where the action is: the foundations of embodied interaction*. MIT press.

[25] Stefania Druga, Sarah T.Vu, Eesh Likhith, and Tammy Qiu. 2019. Inclusive AI literacy for kids around the world.

[26] Allison Druin, Benjamin B Bederson, Juan Pablo Hourcade, Lisa Sherman, Glenda Revelle, Michele Platner, and Stacy Weng. 2001. Designing a digital library for young children. In *Proceedings of the 1st ACM/IEEE-CS joint conference on Digital libraries*, 398–405.

[27] Madeleine Clare Elish. 2019. Moral crumple zones: Cautionary tales in human-robot interaction. *Engaging Science, Technology, and Society* 5: 40–60.
[28] Noel Enyedy, Joshua A Danish, Girlie Delacruz, and Melissa Kumar. 2012. Learning physics through play in an augmented reality environment. *International journal of computer-supported collaborative learning* 7, 3: 347–378.
[29] John H Falk and Lynn D Dierking. 2000. *Learning from museums: Visitor experiences and the making of meaning*. Altamira Press.
[30] William R Frey, Desmond U Patton, Michael B Gaskell, and Kyle A McGregor. 2020. Artificial intelligence and inclusion: Formerly gang-involved youth as domain experts for analyzing unstructured Twitter data. *Social Science Computer Review* 38, 1: 42–56.
[31] Helen L Gallagher and Christopher D Frith. 2003. Functional imaging of 'theory of mind.' *Trends in cognitive sciences* 7, 2: 77–83.
[32] William Gaver. 2011. Making spaces: how design workbooks work. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 1551–1560.
[33] William Gaver. 2012. What should we expect from research through design. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 937–946.
[34] Erving Goffman. 2008. *Behavior in public places*. Simon and Schuster.
[35] Norma González, Luis C Moll, and Cathy Amanti. 2006. *Funds of knowledge: Theorizing practices in households, communities, and classrooms*. Routledge.
[36] Google. 2021. Experiments with Google. *AI Experiments*. Retrieved January 7, 2020 from https://experiments.withgoogle.com/collection/ai
[37] David Gunning. 2017. Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA), nd Web*.
[38] Joshua P Gutwill and Sue Allen. 2010. Facilitating family group inquiry at science museum exhibits. *Science Education* 94, 4: 710–742.
[39] Mark Guzdial. 2013. Exploring Hypotheses about Media Computation. In *Proceedings of the Ninth Annual International ACM Conference on International Computing Education Research*, 19–26.
[40] Christian Heath and Dirk vom Lehn. 2008. Configuring "Interactivity" Enhancing Engagement in Science Centres and Museums. *Social Studies of Science* 38, 1: 63–91.
[41] Rafael Lozano Hemmer. 2015. *Level of Confidence*. Retrieved from http://www.lozano-hemmer.com/artworks/level_of_confidence.php
[42] Suzanne Hidi and K Ann Renninger. 2006. The four-phase model of interest development. *Educational psychologist* 41, 2: 111–127. https://doi.org/10.1207/s15326985ep4102_4
[43] Jon Hindmarsh, Christian Heath, Dirk Vom Lehn, and Jason Cleverly. 2005. Creating assemblies in public environments: Social interaction, interactive exhibits and CSCW. *Computer Supported Cooperative Work (CSCW)* 14, 1: 1–41.
[44] Michael S Horn, Erin Treacy Solovey, R Jordan Crouser, and Robert JK Jacob. 2009. Comparing the use of tangible and graphical programming languages for informal science education. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 975–984.
[45] Michael S. Horn, Erin Treacy Solovey, and Robert JK Jacob. 2008. Tangible programming and informal science learning: making TUIs work for museums. In *Proceedings of the 7th international conference on Interaction design and children*, 194–201. Retrieved October 31, 2015 from http://dl.acm.org/citation.cfm?id=1463756
[46] Eva Hornecker. 2005. A design theme for tangible interaction: embodied facilitation. In *ECSCW 2005*, 23–43.
[47] Eva Hornecker and Jacob Buur. 2006. Getting a grip on tangible interaction: a framework on physical space and social interaction. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 437–446.
[48] Eva Hornecker and Luigina Ciolfi. 2019. Human-computer interactions in museums. *Synthesis Lectures on Human-Centered Informatics* 12, 2: i–171.
[49] Thomas Humphrey, Joshua Gutwill, and The Exploratorium APE Team. 2005. *Fostering Active Prolonged Engagement: The Art of Creating APE Exhibits*. Routledge, Abingdon, UK.

[50] Mizuko Ito, Kris Gutiérrez, Sonia Livingstone, Bill Penuel, Jean Rhodes, Katie Salen, Juliet Schor, Julian Sefton-Green, and S Craig Watkins. 2013. *Connected learning: An agenda for research and design*. Digital Media and Learning Research Hub.
[51] Scott R Klemmer, Björn Hartmann, and Leila Takayama. 2006. How bodies matter: five themes for interaction design. In *Proceedings of the 6th conference on Designing Interactive systems*, 140–149.
[52] Sheila Krogh-Jespersen, Kimberly A Quinn, William Leo Donald Krenzer, Christine Thi Nguyen, Jana Greenslit, and Aaron Price. 2020. Exploring the Awe-some: Mobile eye-tracking insights into awe in a science museum.
[53] George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
[54] Jean Lave and Etienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge, UK.
[55] Annabel Lindner, Stefan Seegerer, and Ralf Romeike. 2019. Unplugged Activities in the Context of AI. In *International Conference on Informatics in Schools: Situation, Evolution, and Perspectives*, 123–135.
[56] Lucas Liu, Duri Long, and Brian Magerko. 2020. MoViz: A Visualization Tool for Comparing Motion Capture Data Clustering Algorithms. In *Proceedings of the 7th International Conference on Movement and Computing*, 1–8.
[57] Daria Loi, Thomas Lodato, Christine T Wolf, Raphael Arar, and Jeanette Blomberg. 2018. PD manifesto for AI futures. In *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial-Volume 2*, 1–4.
[58] Duri Long, Mikhail Jacob, Nicholas Davis, and Brian Magerko. 2017. Designing for Socially Interactive Systems. In *Proceedings of the 11th Conference on Creativity and Cognition*.
[59] Duri Long, Mikhail Jacob, and Brian Magerko. 2019. Designing Co-Creative AI for Public Spaces.
[60] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 ACM Conference on Human Factors in Computing Systems (CHI 2020)*. https://doi.org/10.1145/3313831.3376727
[61] Duri Long, Tom McKlin, Anna Weisling, William Martin, Steven Blough, Katlyn Voravong, and Brian Magerko. 2020. Out of Tune: Discord and Learning in a Music Programming Museum Exhibit. In *Proceedings of the 2020 ACM Conference on Interaction Design and Children (IDC'20)*. https://doi.org/10.1145/3392063.3394430
[62] Duri Long, Tom McKlin, Anna Weisling, William Martin, Hannah Guthrie, and Brian Magerko. 2019. Trajectories of Physical Engagement and Expression in a Co-Creative Museum Installation. In *Proceedings of the 2019 ACM Conference on Creativity and Cognition*. https://doi.org/10.1145/3325480.3325505
[63] Martin Ludvigsen. 2005. Designing for social use in public places–A conceptual framework of social interaction. *Proceedings of Designing Pleasurable Products and Interfaces, DPPI* 5: 389–408.
[64] Caitlin Lustig. 2019. Intersecting Imaginaries: Visions of Decentralized Autonomous Systems. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW: 1–27.
[65] Brian Magerko, Jason Freeman, Tom McKlin, Mike Reilly, Elise Livingston, Scott McCoid, and Andrea Crews-Brown. 2016. EarSketch: A STEAM-Based Approach for Underrepresented Populations in High School Computer Science Education. *ACM Transactions on Computing Education (TOCE)* 16, 4: 14. https://doi.org/10.1145/2886418
[66] Ann Mintz. 2005. Science, society and science centres.
[67] Museum of Science and Industry, Chicago. 2018. *Robot Revolution*. Retrieved from https://www.msichicago.org/explore/whats-here/exhibits/robot-revolution/
[68] Tomohiro Nishida, Susumu Kanemune, Yukio Idosaka, Mitaro Namiki, Tim Bell, and Yasushi Kuno. 2009. A CS unplugged design pattern. *ACM SIGCSE Bulletin* 41, 1: 231–235.
[69] Lorelli S Nowell, Jill M Norris, Deborah E White, and Nancy J Moules. 2017. Thematic analysis: Striving to meet the trustworthiness criteria. *International Journal of Qualitative Methods* 16, 1: 1609406917733847.
[70] Hyunjoo Oh, Anisha Deshmane, Feiran Li, Ji Yeon Han, Matt Stewart, Michael Tsai, Xing Xu, and Ian Oakley. 2013. The digital dream lab: tabletop puzzle blocks for exploring programmatic concepts. In *Proceedings of the 7th international conference on tangible, embedded and embodied interaction*, 51–56.
[71] Marina Papastergiou. 2008. Are computer science and information technology still masculine fields? High school students' perceptions and career choices. *Computers and Education* 51, 2: 594–608.

[72] Eli Pariser. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
[73] Aaron Price and B Pernot. 2015. *Educational & Demographic Backgrounds of Guests on the Museum Floor*. Museum of Science and Industry, Chicago.
[74] ReadyAI. 2020. *Cozmo as a Self-Driving Car Exhibit at Carnegie Science Center*. Retrieved from https://www.youtube.com/watch?v=1XygLnR1Gd8
[75] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Kafai, Yasmin. 2009. Scratch: Programming for All. *Communications of the ACM* 52, 11: 60–67. https://doi.org/10.1145/1592761.1592779
[76] Jessica Roberts and Leilah Lyons. 2017. Scoring Qualitative Informal Learning Dialogue: The SQuILD Method for Measuring Museum Learning Talk. . Philadelphia, PA: International Society of the Learning Sciences.
[77] Kasper Rodil, Matthias Rehm, and Antonia Lina Krummheuer. 2018. Co-designing social robots with cognitively impaired citizens. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*, 686–690.
[78] Liz Sanders. 2012. Convivial Toolbox: Generative Research for the Front End of Design pdf by.
[79] Joe Saunders, Dag Sverre Syrdal, Kheng Lee Koay, Nathan Burke, and Kerstin Dautenhahn. 2015. "Teach Me–Show Me"—end-user personalization of a smart home and companion robot. *IEEE Transactions on Human-Machine Systems* 46, 1: 27–40.
[80] Ben Rydal Shapiro, Rogers P Hall, and David A Owens. 2017. Developing & using interaction geography in a museum. *International Journal of Computer-Supported Collaborative Learning* 12, 4: 377–399.
[81] Zoe Skinner, Stacey Brown, and Greg Walsh. 2020. Children of Color's Perceptions of Fairness in AI: An Exploration of Equitable and Inclusive Co-Design. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–8.
[82] Matthew Smith, Christian Szongott, Benjamin Henne, and Gabriele Von Voigt. 2012. Big data privacy issues in public social media. In *2012 6th IEEE International Conference on Digital Ecosystems and Technologies (DEST)*, 1–6.
[83] Scott S Snibbe and Hayes S Raffle. 2009. Social immersive media: pursuing best practices for multi-user interactive camera/projector exhibits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1447–1456.
[84] Karen Sullenger. 2006. Beyond School Walls: Informal Education and the Culture of Science. *Education Canada* 46, 3: 15–18.
[85] Elisabeth Sulmont, Elizabeth Patitsas, and Jeremy R Cooperstock. 2019. Can You Teach Me To Machine Learn? In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, 948–954.
[86] The Barbican. 2019. *AI: More Than Human*. Retrieved from https://www.barbican.org.uk/whats-on/2019/event/ai-more-than-human
[87] Peter Tolmie, Steve Benford, Chris Greenhalgh, Tom Rodden, and Stuart Reeves. 2014. Supporting group interactions in museum visiting. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, 1049–1059.
[88] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What should every child know about AI? In *Proceedings of the 2019 Conference on Artificial Intelligence*.
[89] David S Touretzky. 2017. Computational thinking and mental models: From kodu to calypso. In *Blocks and Beyond Workshop (B&B), 2017 IEEE*, 71–78.
[90] Lev Semenovich Vygotsky. 1980. *Mind in society: The development of higher psychological processes*. Harvard university press.
[91] Greg Walsh, Alison Druin, Mona Leigh Guha, Elizabeth Foss, Evan Golub, Leshell Hatley, Elizabeth Bonsignore, and Sonia Franckel. 2010. Layered elaboration: a new technique for co-design with children. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1237–1240.
[92] Greg Walsh and Elizabeth Foss. 2015. A case for intergenerational distributed co-design: the online kidsteam example. In *Proceedings of the 14th International Conference on Interaction Design and Children*, 99–108.

[93] Danding Wang, Qian Yang, Ashraf Abdul, and Brian Y Lim. 2019. Designing theory-driven user-centric explainable AI. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1–15.
[94] Henry M Wellman, David Cross, and Julanne Watson. 2001. Meta-analysis of theory-of-mind development: The truth about false belief. *Child development* 72, 3: 655–684.
[95] Linda L. Werner, Brian Hanks, and Charlie McDowell. 2004. Pair-programming helps female computer science students. *Journal on Educational Resources in Computing (JERIC)* 4, 1: 4.
[96] Niels Wouters, John Downs, Mitchell Harrop, Travis Cox, Eduardo Oliveira, Sarah Webber, Frank Vetere, and Andrew Vande Moere. 2016. Uncovering the honeypot effect: How audiences engage with public interactive systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 5–16.
[97] Jason C Yip, Tamara Clegg, June Ahn, Judith Odili Uchidiuno, Elizabeth Bonsignore, Austin Beck, Daniel Pauw, and Kelly Mills. 2016. The evolution of engagements and social bonds during child-parent co-design. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 3607–3619.
[98] Xiaofei Zhou, Jessica Van Brummelen, and Phoebe Lin. 2020. Designing AI Learning Experiences for K-12: Emerging Works, Future Opportunities and a Design Framework. *arXiv preprint arXiv:2009.10228*.
[99] Heather Toomey Zimmerman, Suzanne Reeve, and Philip Bell. 2010. Family sense-making practices in science center conversations. *Science Education* 94, 3: 478–505.
[100] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 493–502.
[101] Abigail Zimmermann-Niefield, Makenna Turner, Bridget Murphy, Shaun K Kane, and R Benjamin Shapiro. 2019. Youth Learning Machine Learning through Building Models of Athletic Moves. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*, 121–132.