

Designing for Public Enlightenment: Enhancing Generative AI Literacy on Socio-technical Aspects in Informal Learning Spaces

Chengzhi Zhang
czhang694@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

ABSTRACT

Generative AI presents a critical paradox: widespread public use coexists with a significant “literacy gap” in understanding its true socio-technical perils and benefits, including misinformation, bias, and environmental costs. This research addresses this gap by developing and evaluating interactive learning interventions for informal settings, such as museums, libraries, and parks. Through a mixed-methods approach, the research will identify public misconceptions, develop a pedagogical framework for AI literacy, and create tangible and embodied exhibits to foster more informed public engagement with generative AI. The goal is to empower users to engage responsibly with generative AI, moving beyond superficial use to a more democratic and informed approach.

CCS CONCEPTS

• **Applied computing** → *Interactive learning environments*; • **Social and professional topics** → **Computing education**; • **Human-centered computing** → *Interaction design*.

KEYWORDS

Generative AI literacy, Interaction Design, Embodied Interaction

ACM Reference Format:

Chengzhi Zhang. 2018. Designing for Public Enlightenment: Enhancing Generative AI Literacy on Socio-technical Aspects in Informal Learning Spaces. In *Proceedings of Tangible, Embedded, and Embodied Interaction (TEI '26)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/XXXXXXX.XXXXXX>

1 INTRODUCTION

Generative Artificial Intelligence (generative AI, broadly referring to large-scale, pretrained generative models, e.g., ChatGPT¹, Deepseek², Gemini³) has gained worldwide popularity over the past several years. This interest surged into a global phenomenon following the launch of ChatGPT-3.5 in November 2022 [19]. It has also presented a profound duality, offering immense societal

benefits and systemic risks. While the use of tools like ChatGPT is phenomenal, a significant gap persists between the public’s surface-level understanding and a deeper “generative AI literacy”—the ability to critically evaluate the technology’s complex ethical, social, and economic implications [4]. Those nuances are usually not captured in the widely-adopted general AI literacy guidelines [16, 21]. This “literacy gap” is often fueled by pervasive, often misleading, cultural narratives and over-hyped advertisements that can distort public discourse and hinder a decent understanding. The gap can lead to severe consequences, including harming professionalism in the workplace [20], compromising integrity in academic settings [2], spiking disinformation and misinformation spread [27], and introducing fatal consequences to adults and minors [5, 11, 24]. The research addresses this gap by focusing on the potential of informal learning spaces (e.g., museums, libraries, maker spaces) as target sites for public education. We target the general public as subjects, including both adults and minors aged 12 and above. This age threshold aligns with the onset of Piaget’s formal operational stage [14], a period during which abstract thought and hypothetical reasoning develop, and is essential for cultivating AI literacy. This research is structured around three key sequential Research Questions (RQs):

- RQ1: What are the general public’s perceived socio-technical perils and benefits towards generative AI? What are the lingering socio-technical conceptions and misconceptions towards generative AI?
- RQ2: What are the necessary learning objectives to convey to the public to debunk these misconceptions and inappropriate mental models, and foster their generative AI literacy, based on our findings of RQ1?
- RQ3: How do we design learning interventions in informal learning spaces focusing on tangible and embodied interactions to convey the learning objectives specified in RQ2?

Through a mixed-methods approach combining user studies (interviews and questionnaires) and design-based research, the research plan is to produce three key contributions:

- An empirical study results on public conceptions and misconceptions of generative AI;
- A pedagogical framework of core learning objectives for public generative AI literacy (with a focus on social-technical, not on pure technical aspects);
- A set of evaluated, interactive interventions deployed in informal learning environments.

This research aims to empower more informed public engagement and democratic participation in the governance of generative AI. To that end, the following sections outline research stages, methods, and measurable outputs:

¹ChatGPT: <https://chatgpt.com/>

²Deepseek: <https://www.deepseek.com/>

³Gemini: <https://gemini.google.com/app>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

TEI '26, March 8-11, 2026, Chicago, Illinois, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/XXXXXXX.XXXXXX>

2 UNDERSTANDING PUBLIC CONCEPTIONS AND MISCONCEPTIONS

- Research method(s): Questionnaire.
- Measurable output(s): A detailed user study report presenting both qualitative themes and quantitative results, providing a rich, evidence-based answer to RQ1.

This foundational research phase will establish an empirically grounded understanding of the public’s current conceptions and misconceptions of generative AI. Many existing common misconceptions are derived from science fiction and over-hyped advertisements, such as the belief that AI is sentient, infallible, or will inevitably “take over the world” [7, 8]. Beyond these existing myths about AI, we plan to employ an explanatory mixed-methods approach to unravel misconceptions of generative AI specifically.

First, we will review current research and synthesize existing knowledge on documented socio-technical impacts—including benefits such as enhanced productivity and creativity [10], and risks such as job displacement [1], algorithmic bias [6], environmental costs [9], and misinformation [27]. The review will inform the first user study, which will include a questionnaire comprising both closed-ended questions (for quantitative data) and open-ended questions (for qualitative data) that explore perceived perils, benefits, and the nuanced narratives from the public. The questionnaire will be publicly distributed and will quantitatively measure the prevalence of the specific conceptions and misconceptions identified, allowing for a robust triangulation of data. Through reviewing existing research and questionnaires, we aim to form a solid understanding of the public’s conceptualization of the socio-technical perils and benefits of generative AI.

3 CONSTRUCTING LEARNING OBJECTIVES

- Research method(s): Quantitative and qualitative data analysis.
- Measurable output(s): A document on “Learning Guidelines and Defined Learning Objectives” presenting the pedagogical framework that specifies clear competency areas.

This research phase will serve as the critical bridge between the diagnostic findings of RQ1 and the system contributions of RQ3, translating the empirically identified public misconceptions and knowledge gaps into a structured, prioritized, and actionable pedagogical framework. Our early-stage scoping review has produced a framework for generative AI literacy, which will inform RQ2 [26]. The key themes, knowledge gaps, and prevalent misconceptions identified in the user study report from RQ1 will be distilled into a concise list of “learning objectives.” For example, a finding that “a majority of participants are unaware of the significant power consumption of generative AI usage” translates to a solid learning need regarding generative AI’s environmental impact. The results of the questionnaire from RQ1 will be analyzed, with methods including thematic analysis for open-ended questions [23].

Finally, the learning needs will be systematically mapped onto learning objectives, allowing for prioritization based on the prevalence of the misconception and the importance in expert discourse. From this mapping, we will formulate specific and measurable learning objectives, such as “learners will be able to realize the bias in generative AI results” to counter misplaced confidence about the

objective results. Overall, the RQ2 will produce a document listing “Learning Guidelines and Defined Learning Objectives” that can serve not only our subsequent intervention design, but can also inform relevant research on generative AI literacy and generative AI systems designs.

4 IMPLEMENTING LEARNING INTERVENTIONS AND EVALUATION

- Research method(s): System design and implementation, user evaluation (pre-/post- questionnaire, semi-structured interview, learning talk analysis).
- Measurable output(s): The primary contribution is the interactive systems, which are public-facing and will be displayed on public sites. This will be accompanied by a comprehensive user evaluation report detailing the effectiveness of each intervention and the design principles derived from the iterative design process.

We will conduct a series of brainstorming sessions and workshops to generate concepts for designs that convey the learning objectives concluded from RQ2. Subsequently, we will develop interactive systems based on the selected concepts, adhering to the specified criteria (e.g., ease of conveyance and the importance of the learning objective). The designs will be grounded in established principles of informal learning (e.g., allowing for self-directed exploration, and providing multiple entry points) [13, 15]. Tangible and embodied interaction has the merits of enhanced learner engagement, lowering interaction thresholds, and improved learning outcomes [18], which will be the focus of our interaction designs. This might include an interactive exhibit where “asking ChatGPT a question” causes a container of water to visibly drain and a meter measuring electricity use, connecting the seemingly immaterial act of a query to a tangible resource cost.

Until now, we have successfully deployed three interactive installations. *Fool Your Friend* (see Fig. 1 left) is an interactive experience where the learner engages in guessing which post is AI-generated among several posts, and generating their own “fooling” post, teaching that generative AI can produce misinformation and disinformation, and learners need to read the results with skepticism. *Chatbot of Truth* (see Fig. 1 middle) is an interactive installation where the learner asks a human-height robotic head a yes-or-no question, which teaches that chatbots (like ChatGPT) are suitable for certain question types and not suitable for others. *LuminAI* [25] (see Fig. 1 right) is a three-panel exhibit fostering AI literacy through open-ended expressive dance with an AI agent. We have conducted studies with approximately 20 participants per exhibit, utilizing a retrospective pre- and post-test Likert scale to assess the knowledge gain. The results showed that both *Chatbot of Truth* and *Fool Your Friend* participants demonstrated increased understanding of specified learning objectives after interaction, with statistical significance. The experience will further inform our future system implementation.

To evaluate the learning outcomes of our future exhibits, short questionnaires administered after interacting with the installation will measure knowledge gain, attitudes related to the specific learning objective, and learners’ interests. Audio recording of the interaction processes, and the learning talk analysis will provide



Figure 1: Learners interacting with *Fool Your Friend* at the Museum of Design Atlanta (Left). Learner interacting with *Chatbot of Truth* in our lab (Middle, with texts of the screen indicating the category of the learner’s question is “General Knowledge and Fact-Checking” and that “Chatbots can share general knowledge in areas like history, science, and technology, but sometimes they can go wrong or make things up. Please double-check for important facts.” Due to the flickering of the TV, the camera could not capture the text on the screen clearly). Learner interacting with *LuminAI* at the Griffin Museum of Science and Industry (Right, with texts on screen “Congrats! We are dance buddy now... Let’s groove and improvise!”)

valuable insight into learners’ reasoning processes, points of confusion, and inform exhibit iterations [3, 17, 22]. Video recordings of how learners engage with the exhibit (e.g., interaction duration and sessions) will provide insights into learners’ engagement [12]. Post-interaction interviews will gather qualitative feedback on the experience, learning outcomes, and direct our exhibit iterations.

5 RESEARCH TIMELINE

My primary advisor is Professor Brian Magerko at Georgia Institute of Technology. For the summer 2026 semester, I plan to review current research and form a solid understanding of learners’ perceptions of generative AI. For the fall 2026 semester, based on our understanding of the public’s perceptions and misconceptions of generative AI, we will begin developing the learning interventions. For the spring 2027 semester, we will exhibit these learning interventions in public spaces and conduct user studies at our museum partners, including the Griffin Museum of Science and Industry in Illinois, the High Museum and the Museum of Design Atlanta in Georgia.

REFERENCES

- [1] Syed AbuMusab. 2024. Generative AI and human labor: who is replaceable? *AI & SOCIETY* 39, 6 (2024), 3051–3053.
- [2] Hussam Alkaiisi and Samy I McFarlane. 2023. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus* 15, 2 (2023).
- [3] Sue Allen. 2003. Looking for learning in visitor talk: A methodological exploration. In *Learning conversations in museums*. Routledge, 265–309.
- [4] Ravinithesh Annappureddy, Alessandro Fornaroli, and Daniel Gatica-Perez. 2025. Generative AI literacy: Twelve defining competencies. *Digital Government: Research and Practice* 6, 1 (2025), 1–21.
- [5] Imane El Atillah. 2023. Ai chatbot blamed for “encouraging” young father to take his own life. <https://www.euronews.com/next/2023/03/31/man-ends-his-life-after-an-ai-chatbot-encouraged-him-to-sacrifice-himself-to-stop-climate->
- [6] Joan Palmiter Bajorek. 2019. Voice recognition still has significant race and gender biases. *Harvard Business Review* 10 (2019), 1–4.
- [7] Yasmine Belghith, Atefeh Mahdavi Goloujeh, Brian Magerko, Duri Long, Tom Mcklin, and Jessica Roberts. 2024. Testing, Socializing, Exploring: Characterizing Middle Schoolers’ Approaches to and Conceptions of ChatGPT. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [8] Arne Bewersdorff, Xiaoming Zhai, Jessica Roberts, and Claudia Nerdel. 2023. Myths, mis- and preconceptions of artificial intelligence: A review of the literature. *Computers and Education: Artificial Intelligence* 4 (2023), 100143.
- [9] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- [10] Anil R Doshi and Oliver P Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science advances* 10, 28 (2024), eadn5290.
- [11] Clare Duffy. 2025. Parents of 16-year-old sue openai, claiming Chatgpt advised on his suicide | CNN business. <https://www.cnn.com/2025/08/26/tech/openai-chatgpt-teen-suicide-lawsuit>
- [12] Eva Hornecker and Luigina Ciolfi. 2022. *Human-computer interactions in museums*. Springer Nature.
- [13] Thomas Humphrey, Joshua Gutwill, Exploratorium APE Team, et al. 2017. *Fostering active prolonged engagement*. Taylor & Francis New York.
- [14] Deanna Kuhn. 1979. The significance of Piaget’s formal operations stage in education. *Journal of Education* 161, 1 (1979), 34–50.
- [15] Duri Long, Takeria Blunt, and Brian Magerko. 2021. Co-designing AI literacy exhibits for informal learning spaces. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–35.
- [16] Duri Long and Brian Magerko. [n. d.]. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2020) (*CHI ’20*). Association for Computing Machinery, 1–16. <https://doi.org/10.1145/3313831.3376727>
- [17] Duri Long, Anthony Teachey, and Brian Magerko. 2022. Family learning talk in AI literacy learning activities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [18] Paul Marshall. 2007. Do tangible interfaces enhance learning?. In *Proceedings of the 1st international conference on Tangible and embedded interaction*. 163–170.
- [19] Steve Mollman. 2022. CHATGPT gained 1 million users in under a week. here’s why the AI chatbot is primed to disrupt search as we know it. <https://finance.yahoo.com/news/chatgpt-gained-1-million-followers-224523258.html>
- [20] LARRY NEUMEISTER. 2023. *Lawyers blame ChatGPT for tricking them into citing bogus case law*. <https://apnews.com/article/artificial-intelligence-chatgpt-courts-e15023d7e6fdf4f099aa122437dbb59b>
- [21] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence* 2 (2021), 100041.
- [22] Jessica Roberts and Leilah Lyons. 2017. The value of learning talk: applying a novel dialogue scoring method to inform interaction design in an open-ended, embodied museum exhibit. *International Journal of Computer-Supported Collaborative Learning* 12, 4 (2017), 343–376.
- [23] Gareth Terry, Nikki Hayfield, Victoria Clarke, Virginia Braun, et al. 2017. Thematic analysis. *The SAGE handbook of qualitative research in psychology* 2, 17–37 (2017), 25.
- [24] Chloe Xiang. 2023. *‘He Would Still Be Here’: Man Dies by Suicide After Talking with AI Chatbot, Widow Says*. <https://www.vice.com/en/article/wxjxg/openai-and-microsoft-sued-for-dollar3-billion-over-alleged-chatgpt-privacy-violations>
- [25] Chengzhi Zhang, Chelsi Alise Cocking, Milka Trajkova, Zoe Lacy Mock, Gemma Tate, Cassandra Naomi Monden, and Brian Magerko. 2024. Fostering AI Literacy with LuminAI through Embodiment and Creativity in Informal Learning Spaces.

- In *Proceedings of the 16th Conference on Creativity & Cognition*. 476–481.
- [26] Chengzhi Zhang and Brian Magerko. 2025. Generative AI Literacy: A Comprehensive Framework for Literacy and Responsible Use. *arXiv preprint arXiv:2504.19038* (2025).
- [27] Jiawei Zhou, Yixuan Zhang, Qianni Luo, Andrea G Parker, and Munmun De Choudhury. 2023. Synthetic lies: Understanding ai-generated misinformation and evaluating algorithmic and human solutions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.