

AI Literacy: Finding Common Threads between Education, Design, Policy, and Explainability

Post-Workshop Report

Leadership

The workshop was co-organized by Duri Long (Northwestern University), Jessica Roberts (Georgia Tech), and Brian Magerko (Georgia Tech). The main organizers additionally recruited content experts to assist in the reviewing process and moderate panels at the workshop. Panel moderators included Daniella DiPaola (MIT), Fred Martin (University of Massachusetts Lowell), and Ken Holstein (Carnegie Mellon University). Moderators prepared questions for their panels and moderated discussions between panelists and other workshop participants.

Participation

This was an in-person workshop held at [CHI 2023](#) in Hamburg, Germany on April 28, 2023. We received 33 submissions and accepted a total of 22 submissions. 38 people in total (including organizers) attended the workshop. Submissions were selected by a jury. We aimed to maximize attendance—given space and time constraints—to promote the development of this new community at CHI, while also taking into consideration factors such as submission quality, topical relevance, and representing a diversity of perspectives on AI literacy.

We received funding to support the workshop from the National Science Foundation (DRL 2214463). Funding was used to support participant travel to the workshop, with a particular focus on supporting participants who are from groups underrepresented in STEM, are graduate or undergraduate students, and/or are coming from a community that would not otherwise attend CHI (e.g., educators, museum staff, academics from a field such as public policy). Ultimately, 10 attendees received funding to support their travel to the conference. Of these 10 participants, all identified as either being a member of a group marginalized in computing (7/10) and/or being a student (7/10). Of the three participants who were not students, two were in an academic position with minimal support for travel and one was a museum educator. 7/10 recipients responded “Absolutely” to the prompt: “Would receiving financial support make a significant difference in your ability to participate in the workshop?” Finally, 7/10 recipients were first-time CHI attendees.

Of particular note is that the NSF funding supported the attendance of one member of the exhibit staff at the Museum of Science and Industry, Chicago (MSI). MSI is working in collaboration with the workshop organizers to develop museum exhibits to foster AI literacy. Having a MSI staff member in attendance brought a valuable perspective from the informal science learning community to the workshop and additionally helped to familiarize the museum team with ongoing research on AI literacy.

Structure

The workshop was structured as a series of moderated discussion panels on different topics related to AI literacy. Panels were composed of 5-6 participants, each of whom authored a short paper related to the panel topic. All papers were made available to other participants in advance of the workshop via the [workshop website](#). Panelists did not deliver traditional paper presentations, but instead engaged in discussions led by our panel moderators. At scheduled times between panels, participants were encouraged to engage in an affinity mapping activity, using post-it notes to identify common themes from the panel discussions and working in groups to associate similar themes.

During the lunch break, the organizers assigned lunch groups to encourage workshop participants from different universities to connect with each other and to foster mentoring relationships. Group assignments were formulated in advance of the workshop. Each group included both faculty members and students and participants from a variety of different universities.

Takeaways

In this section, we summarize each of the panel discussions and present the final takeaways from the workshop.



Figure 1. Panelists in the “Designing Learning Experiences” panel (left). Participants engaging in the concept mapping activity (right).

Panel 1: Situating AI Literacy

The *Situating AI Literacy* panel was moderated by Duri Long. The submissions in this panel are listed below, with the participating authors bolded.

- *Boosting AI literacy with Explainable AI*
Matija Franklin (University College London) and Stefan Herzog (Max Planck Institute for Human Development)
- *Applying Interdisciplinary Frameworks to Understand Algorithmic Decision-Making*
Timothée Schmude, Laura Koesten, Torsten Möller, Sebastian Tschatschek (University of Vienna, Austria)
- *Exploring AI Literacy in SMEs*
Maike Harbers and Anja Overdiek (Rotterdam University of Applied Sciences)
- *Data Dreams: Raising Collective Awareness Through Literacy Workshops*
Giovanna Nunes Vilazza, Christoffer Bagger, Stine Lomborg (University of Copenhagen)
- *Game Jamming for AI Literacy*
Jeannette Falk (University of Salzburg)
- *AiLingo – Advancing AI Literacy through a Gamified Learning App*
Marc Pinski, **Miguel-José Haas**, and Alexander Benlian (Technical University of Darmstadt)

This panel focused on how to situate AI literacy interventions in varying contexts—including game jams, enterprise workplaces, community workshops, and informal learning spaces for adults. Two papers in this session discussed how AI literacy relates to explainable AI (XAI) and how we can incorporate learning in XAI tools. Throughout the session, participants emphasized the importance of providing opportunities for different audiences to broaden their AI literacy due AI’s wide-ranging impact on areas such as the democratic process, creative labor, and enterprise workplaces.

Given the varying contexts, panelists discussed how to define AI literacy. The group engaged in conversation about whether AI literacy should be defined from the bottom-up (e.g., via interviews with stakeholders) or top-down (e.g., via expert opinions), whether AI literacy is a general-purpose skillset or differs depending on context, and how to prioritize which elements of AI literacy to convey given limited time. A workshop attendee questioned whether the term *literacy* was a help or a hindrance; terms such as *knowledge* and *experience* were floated by one panelist as possible alternatives.

Factors such as learners’ age, occupation, socioeconomic status, and learning environments may impact how AI literacy is defined and communicated. Participants also emphasized that different framings of AI may be needed for different groups—for example, understanding how AI impacts finances may be interesting for adults in an enterprise workplace, but is less relevant in a third-grade classroom. It also may be important to teach adults how to apply AI in specific contexts, whereas children may be building more general-purpose knowledge.

The group also discussed different modalities for evaluating AI literacy, such as explaining design decisions made in a game jam, administering a validated survey instrument as a pre/post test, and assessing semantic (i.e., terminology) vs. procedural (i.e., how to use AI) knowledge.

The discussion surrounding the XAI papers in particular raised questions surrounding how to build trust in an era of misinformation. As the development of XAI shifts away from tools for developers to tools for the public, it is unclear whether many people will even trust the explanations being offered by an AI. Approaches discussed for making AI systems more interpretable for non-experts included reducing the amount of information presented to the user and providing opportunities for interactivity.

Panel 2: AI in K-12 Education

The *AI in K-12 Education* panel was moderated by Fred Martin. The submissions in this panel are listed below, with the participating authors bolded.

- *Towards an Inclusive AI Education: Overcoming Challenges and Promoting Equity in K12 AI Curriculum Development*
Gianluca Schiavo and Valeria Fabretti (Bruno Kessler Foundation)
- *CRAFT-work: An Integrative Co-Design Approach for Designing High School AI Literacy Resources*
Victor R. Lee, Parth Sarin, Jacob Wolf, and Benjamin Xie (Stanford University)
- *Viewing the Finnish national curriculum through AI glasses: participatory design for integrating AI in grades 1-9*
Linda Mannila (Linköping University)
- *Teacher Learning for Student Learning: Prioritizing K-12 Educators’ AI Literacy*
Daniela Ganelin (Stanford University)
- *Everyday AI: Promoting Teacher AI Literacy and Communities of Practice through Teacher Professional Development*
Kate Moore (MIT STEP Lab), Irene Lee (MIT STEP Lab), Helen Zhang (Boston College)

- *Addressing Data Literacy in K-12 Artificial Intelligence Education*
Viktoriya Olari, Kamilla Tenório, Ralf Romeike (Freie Universität Berlin)

This panel focused on the ways in which AI education is being incorporated into K-12 classrooms, including how K-12 teachers are being trained to teach about AI. The panelists brought a variety of international perspectives to the topic, and discussions touched on the variability in political interest in AI curricula, school resources and educational standards, and cultural perspectives on AI across a wide variety of contexts.

Panelists discussed incorporating AI education in programming courses, considering challenges such as how to make the vocabulary and examples more inclusive and what a learning progression might look like. Cross-disciplinary AI education was also discussed by numerous panelists. For instance, one panelist talked about teaching about AI and data literacy through historical data. Many teachers who did not see AI as relevant to their subject areas several years ago are now more open to the idea of incorporating AI education in their classroom. One panelist discussed the importance of leveraging disciplinary teachers' existing expertise for AI education—for example, drawing on a social studies teacher's skills in engaging students in challenging conversations about social or a mathematics teacher's ability to engage students in rigorous analytical thinking.

Numerous approaches for teacher professional development were discussed. Multiple participants emphasized the importance of involving teachers in the design of learning materials for the classroom. Teachers are a unique population, as they need to both develop knowledge about AI but also feel a sense of mastery so that they can explain it to others. This requires both AI knowledge and pedagogical content knowledge about how to teach AI to their students. Teachers' values towards AI—such as their anxiety about AI, perceptions of AI, or personal self-efficacy in AI—play a role in their teaching.

Panelists emphasized the importance of “situated AI” or considering how AI is used in the real world and how it impacts specific students, communities, and families. Teachers need to bring their knowledge of their students to the learning experience and building connections in AI education across classrooms and communities is important.

K-12 AI education was discussed in relation to several other related topics, including data literacy and computer science education. Several groups took lessons directly from earlier efforts to incorporate CS education in classrooms (particularly in a US context). Lessons from professional development for CS teachers included that teachers need to see instruction modeled (rather than just being given resources), they need a community of practice (e.g., a discussion forum, an “AI Book Club”), and we need to distribute training hours across multiple sessions to reduce the cognitive load.

Challenges to incorporating AI in schools were also discussed, and these again drew on lessons from initiatives to incorporate CS education in K-12 schools. Finding time in already packed curricula and political resistance to efforts to modify existing curricula are key issues. Panelists suggested that we need to embrace the idea that teachers are going to pick and choose aspects of an AI curriculum, and that resources should be designed modularly so that teachers can “forage and curate” as needed. This was called a “nooks and crannies” model—for example, a social studies teacher could incorporate an AI activity into a lesson, or a CS teacher who may not have time to teach a full unit on AI could work AI into a one-day class period that they have available between units.

Inclusivity and *access* were key themes that came up throughout the panel—including considering what schools currently have the capacity to teach AI courses, whether teachers have the necessary background

to teach those courses, and which students are going to those select schools. Teaching AI across disciplines fitting it into “nooks and crannies” was one approach suggested to broaden access. Inclusive examples and contexts for introducing AI is also important for reaching students who may not have an interest in or self-efficacy towards the subject. Finally, teaching AI in a social context could help students build connections between their own life experiences and AI. Providing flexible learning materials that teachers can adapt when things do not work as planned is also important—e.g., having contingency plans for students who do not know how to upload files, or for when cameras on school computers are not working.

A final question raised was the role AI could play in helping teachers do their job more efficiently, and the need for both teachers and researchers to have awareness about the varying ways in which AI technologies are impacting education now and in the future.

Panel 3: Designing Learning Experiences

The *Designing Learning Experiences* panel was moderated by Ken Holstein. The submissions in this panel are listed below, with the participating authors bolded.

- *Incorporating AI in a Digital Well-Being Workshop for Middle-School Girls*
Catherine Grevet Delcourt (Wellesley College), Linda Charmaraman (Wellesley Centers for Women)
- *Student-centered design for student-centered learning*
Mary Cate Gustafson-Quiett, Sarah Wharton (MIT)
- *AI & Data Literacy for Non-Technical Students: A Hybrid-Augmented Learning Factory*
Shi Liu, Thimo Schulz, Peyman Toreini, Christian Peukert, Alexander Maedche, Christof Weinhardt (Karlsruhe Institute of Technology)
- *Addressing Artificial Intelligence for Non-computer Science Students in Tertiary Education: Directions for Future Research*
Kamilla Tenório, Viktoriya Olari, Ralf Romeike (Freie Universität Berlin)
- *Towards “Anytime, Anywhere” Community Learning and Engagement around the Design of Public Sector AI*
Wesley Hanwen Deng, Motahhare Eslami, Kenneth Holstein (Carnegie Mellon University)

In this panel, participants discussed the design of learning environments to foster AI literacy across a wide variety of contexts and modalities—including a remote workshop for middle-school girls, classroom and after-school activities for students, a mixed reality tool to support embodied learning about AI, education for non-CS majors at the tertiary level, and learning interventions to support community engagement surrounding public sector AI.

Panelists prioritized communicating different AI literacy learning goals depending on which context they were working in. For example, one panelist discussed AI literacy as the ability for students to build something with AI, moving along a continuum from *unconscious consumer of AI* to *conscious consumer of AI* to *builder of AI*. Another panelist talked about the role of AI education in communities, prioritizing knowledge that enables community members to advocate for and protect themselves when AI is introduced in the public sector. Yet another panelist was interested in exploring how AI education could enable young girls to develop more healthy relationships with social media. Multiple panelists emphasized using co-design to collaboratively decide on learning goals with stakeholders in order to ensure the learning goals serve student interests and relate to their lives.

Multiple modalities for learning were also explored. One panelist discussed designing an immersive learning experience using the Microsoft HoloLens in which students could explore concepts like decision trees in an embodied way. Tradeoffs between modalities also exist—for instance, the HoloLens provides learners with the unique opportunity to step “into the shoes” of an AI by making intangible algorithms tangible, but the technology is expensive and can be inaccessible. Another panelist was forced to conduct remote workshops with students due to COVID, but this format allowed them to reach more young women across a broader geographical range. Multiple panelists emphasized the importance of making connections between the learning interventions and learners’ everyday lives, either via lectures to supplement immersive interactives, or via co-designing learning activities.

Panelists also discussed their hopes for the future of AI literacy research. One panelist was eager to see more research on pre-service training for teachers; another called for more interactive learning experiences; yet another called for more research on empowering everyday users of AI, not just students.

The panelists reflected on possible harms of their work, such as unintentionally contributing to “AI hype.” Although empowering learners to build AI and engage in design decisions about AI can be key learning goals, panelists also did not want to empower learners to build any AI they want without considering ethical ramifications. Concern was also expressed regarding sharing information about the harms of AI with marginalized students without causing additional trauma. A balance needs to be struck, as it is not always empowering to learn about harms when you are powerless to fix them, and it can put the onus of responsibility on students to solve problems they did not create. Panelists concurred that AI education is just one component of more responsible AI development, and that companies and governments also need to take steps to regulate AI technologies.

Finally, panelists reflected on the impact of ChatGPT and other recently developed generative AI technologies on designing AI-related learning experiences. Panelists saw both opportunities and challenges in light of these new tools. Students are starting to see AI as being more clearly relevant to their lives, but at the same time it is difficult to design learning experiences that keep pace with current technology. For example, students learning AI basics by building a chatbot of their own may be disappointed when it doesn’t work as well as a large language model. More research is needed on how to raise awareness of negative impacts of generative AI and to aid learners in leveraging these tools for their own benefit.

Panel 4: AI, Ethics, and Public Policy

The *AI, Ethics, and Public Policy* panel was moderated by Daniella DiPaola. The submissions in this panel are listed below, with the participating authors bolded.

- *Data Hazards*
Vanessa Aisyahsari Hanschke, Natalie Zelenka, Nina Di Cara (University of Bristol)
- *Motivations, Goals, and Pathways for AI Literacy for Journalism*
Mandi Cai*, Sachita Nishal* (Northwestern University) *Co-First Authors
- *Educating “Generation AI” for the cybersecurity mindset*
Jari, Laru (University of Oulu), Matti, Tedre (University of Eastern Finland), Kati Mäkitalo (University of Oulu)
- *How to teach AI without an AI expert in the room? Toward playful, ethical, and accessible AI education*
Andy Stoiber, YJ Kim (UW-Madison)

- *Ethical by Design: Teaching Middle-school Students to Think Ethically About AI*
Glenda S. Stump, Prerna Ravi, Annalisa J. Broski, Angela Daniel, Hal Abelson, Eric Klopfer, Cynthia Breazeal (MIT)

This panel focused on discussing the relationship between public policy, ethics, and AI literacy. Panelists approached this topic from multiple perspectives, including efforts to engage youth in ethical thinking surrounding AI; designing a shared language to identify and reflect on the dangers of AI and data; and considering the role journalists play in shaping AI policy and what AI literacy for journalists looks like.

Multiple panelists touched on the need for shared language and framework for discussing AI ethics. Common frameworks could be used to both scaffold our conversations about AI ethics and spur reflections by developers, ideally leading to changes and reduced harm in technologies. A shared language could also aid in fostering improved communication between those with technical expertise and those with expertise in asking important ethical questions (e.g., social scientists, humanities researchers). The relationship between AI literacy and AI ethics could also be clearer—they are closely intertwined, but one panelist pointed out that people may not need a lot of technical understanding of AI to grasp the consequences of social media companies collecting your data. Training people to ask good questions may be more important in enabling them to advocate for themselves than providing technical expertise.

Several approaches to developing shared languages were discussed. One panelist presented a project in which they were designing “data hazard” labels to identify potential dangers of AI technologies and datasets in a similar way to how we identify chemical hazards. Another panelist discussed the possibility of journalists codifying language / AI terminology in style guides and newsroom practices. Since journalists play a key role in communicating AI technology and policy developments to the public, the ways in which they describe AI and AI ethics are consequential. One group that had been working on AI ethics education for middle school students discussed integrating a framework focused on incorporating diverse perspectives during the technology design process in their classroom activities. Panelists noted that as researchers, it is important for us to consider how we are representing AI in our public-facing presentations, and whether those representations perpetuate misconceptions (e.g., always representing AI with an image of a robot).

However, panelists also acknowledged that frameworks embody value systems. There is a need to listen to multiple perspectives and consider different community and cultural approaches to ethical thinking as we develop shared languages and frameworks for AI ethics. In addition, standardizing frameworks or values can have unintended and even negative consequences—for example, framing removing bias from AI as ‘good’ could imply that there can be a state in which no bias exists, which can be misleading. One panelist drew an analogy to how the scientific method is widely taught and standardized, but many feel that it poorly represents the actual scientific process. Care should be taken to avoid similar harms from AI ethics frameworks.

The discussant raised the question of how to make ethics more central to learning about AI, rather than an add-on that is tagged onto the end of a discussion or lesson. One panelist highlighted the importance of working ethics into all aspects of a curriculum (e.g., by having students role-play as different stakeholders, create speculative fictions about what a future AI might look like in different social contexts, and document their ethical thinking throughout a design activity) rather than tagging on discussion questions at the end of a lesson that teachers often skip over. Another panelist considered how

Panelists working on youth education identified key challenges to engaging students in AI ethics discussions. Teachers may feel unprepared to guide students in a respectful discussion about issues like

racial bias in AI, especially if they are coming from a technical background and do not have training in leading complex discussions. In some cases, teachers may not be allowed to—or may feel scared to—directly engage with certain topics (e.g., in school districts in the U.S. that are actively censoring content about race and gender). Possible solutions discussed included leveraging the expertise of teachers in other disciplines like history, presenting students with the technology and letting them discover its issues on their own (rather than prescribing ‘right’ vs. ‘wrong’), and providing students with a framework to guide their consideration of ethics during design projects.

Finally, policy was discussed as an important lever for spurring industry to develop more ethical and explainable AI technologies. An analogy was made to media literacy – is it enough to tell people to check their sources when numerous sources are now actively trying to deceive users with misinformation? Broadening public AI literacy and ethical literacy can only go so far, and there was a call for more policy regulating usage and/or development of AI technologies.

Concept Mapping

At key points throughout the workshop, we had participants engage in concept mapping with their tables to distill key themes from the panels and engage participants in further discussion about the panel topics. At the end of the workshop, we asked each table to share several key takeaways from their concept mapping discussions. Figure 2 shows an example concept map and the final takeaways.

These concept maps are an important output of this workshop, and to share them back with the community in an accessible format, we have put together both a [document](#) listing the concepts identified by each group (including a list of the final takeaways) and a word cloud that visualizes topic frequency (Figure 3). We have additionally made the concept map photos publicly available on the workshop website.

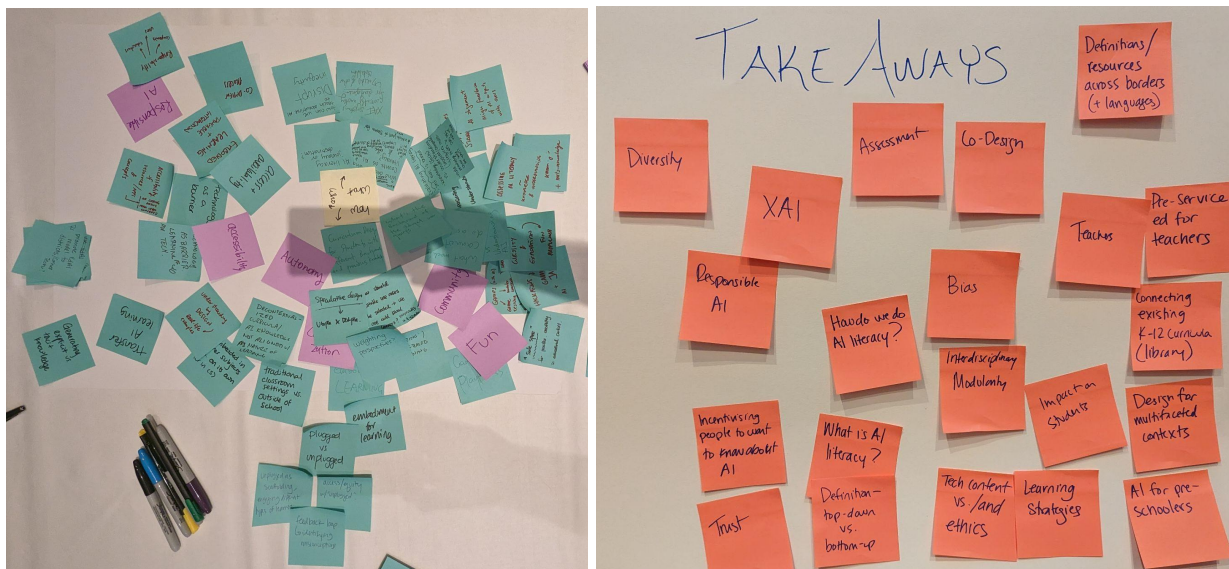


Figure 2. One table’s concept map at the end of the workshop (left). Final takeaways from all of the tables (right).

This workshop helped me to...

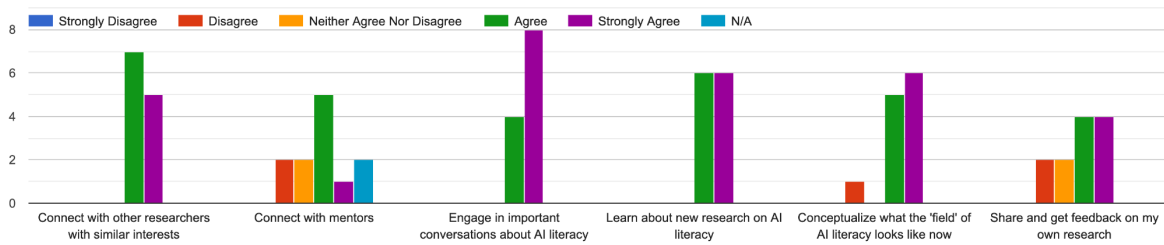


Figure 3. Summary of 12 workshop participants' feedback on various aspects of the workshop.

Resources

We have created a [workshop website](#) which serves as a repository for the workshop outputs. The website can be publicly accessed. The following resources are available on the official workshop website, which will be kept active for at least five years:

- PDFs of all papers presented at the workshop
- Organizer slides from the workshop
- Details about the workshop structure and submission process
- Visual outputs from the concept mapping activity
- This post-workshop report

Future Plans

Based on the strong interest in the workshop this year, the continued interest of the organizers in this topic and community, and enthusiasm from participants, we are planning to propose a second annual AI Literacy workshop at CHI 2024. We envision a similar format that centers on collaborative discussion and ideation and also plan to continue the coordinated mentorship lunches. We may consider novel panel topics and extensions to the concept mapping activity based on submissions received and the evolving discourse surrounding AI literacy.

We would like to involve additional members of the community in organizing the 2024 workshop, both in terms of the core planning and the panel moderation. This would help us to create a sustainable model for future iterations of the workshop. Other improvements we are considering for next year include: publicizing opportunities for financial support in advance of submissions to ensure we reach participants who may not otherwise be able to attend, considering how to support a larger group of attendees given the strong interest in the topic at CHI, and potentially implementing a standardized review process and publishing a more formal proceedings of the workshop, as the state of the field evolves.