



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

Duri Long
Department of Communication
Studies, Northwestern University
duri@northwestern.edu

Jessica Roberts
School of Interactive Computing,
Georgia Tech
jessica.roberts@cc.gatech.edu

Brian Magerko
School of Literature, Media, and
Communication, Georgia Tech
magerko@gatech.edu

Ken Holstein
Human Computer Interaction,
Carnegie Mellon
kenneth.holstein@gmail.com

Daniella Dipaola
Personal Robots Group, MIT Media
Lab
dipaola@mit.edu

Fred Martin
Computer Science, University of
Massachusetts Lowell
fred_martin@uml.edu

ABSTRACT

Fostering public AI literacy has been a growing area of interest at CHI for several years, and a substantial community is forming around issues such as teaching children how to build and program AI systems, designing learning experiences to broaden public understanding of AI, developing explainable AI systems, understanding how novices make sense of AI, and exploring the relationship between public policy, ethics, and AI literacy. Previous workshops related to AI literacy have been held at other conferences (e.g., SIGCSE, AAAI) that have been mostly focused on bringing together researchers and educators interested in AI education in K-12 classroom environments, an important subfield of this area. Our workshop seeks to cast a wider net that encompasses both HCI research related to introducing AI in K-12 education and also HCI research that is concerned with issues of AI literacy more broadly, including adult education, interactions with AI in the workplace, understanding how users make sense of and learn about AI systems, research on developing explainable AI (XAI) for non-expert users, and public policy issues related to AI literacy.

CCS CONCEPTS

• **Social and professional topics** → Professional topics; Computing education; • **Human-centered computing** → Interaction design; • **Computing methodologies** → Artificial intelligence.

KEYWORDS

AI literacy, AI ethics, explainable AI, AI education, AI4K12, public policy

ACM Reference Format:

Duri Long, Jessica Roberts, Brian Magerko, Ken Holstein, Daniella Dipaola, and Fred Martin. 2023. AI Literacy: Finding Common Threads between Education, Design, Policy, and Explainability. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*,

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI EA '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9422-2/23/04.

<https://doi.org/10.1145/3544549.3573808>

April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 6 pages.
<https://doi.org/10.1145/3544549.3573808>

1 BACKGROUND

1.1 Motivation

Fostering public AI literacy (i.e. “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” [1]) has been a growing area of interest at CHI for several years, and a substantial community is forming around issues such as teaching children how to build and program AI systems, designing learning experiences to broaden public understanding of AI, developing explainable AI systems, understanding how novices make sense of AI, and exploring the relationship between public policy and AI literacy. Previous workshops related to AI literacy have been held at other conferences (e.g., SIGCSE, AAAI) that have been mostly focused on bringing together researchers and educators interested in AI education in K-12 classroom environments, an important subfield of this area [2]. Our workshop seeks to cast a wider net that encompasses both HCI research related to introducing AI in K-12 education and also HCI research that is concerned with issues of AI literacy more broadly, including adult education, interactions with AI in the workplace, understanding how users make sense of and learn about AI systems, research on developing explainable AI (XAI) for non-expert users, and public policy issues related to AI education.

There is a demonstrably strong interest in AI literacy within the CHI community, as evidenced by searching CHI Proceedings for terms such as “AI literacy,” “AI education,” and “explainable AI” (125 exact matches, almost all from the past 2-3 years). There is also a lot of AI literacy related research extending beyond CHI into other communities such as AAAI and SIGCSE that we would like to bring into conversation with the AI literacy work happening at CHI. As the topic of AI literacy is a relatively new area of research, the CHI community would benefit from a workshop specifically focused on bringing together researchers and identifying common themes and definitions across research areas.

1.2 Topics for Discussion

In this workshop, we seek to bring together researchers in the CHI community who are conducting research related to one or more of the topics listed below.

- designing formal or informal learning experiences to foster learning about AI (for adults and/or children)
- developing instruments to evaluate learner understanding of AI
- designing tools or curricula to teach about AI
- developing explainable AI (XAI) systems that can foster learning-through-interaction with non-expert human partners
- understanding how non-experts make sense of AI systems and tools
- investigating how novices learn about AI through activities such as information foraging
- designing public art or installations to engage people in learning about AI
- community-based research investigating relevant issues, concerns, and needs related to AI literacy
- developing equitable and culturally responsive AI literacy initiatives
- shaping or studying public policy and ethics related to AI education and AI literacy

We will group workshop submissions into higher-level panels on Explainable AI, Designing Interactive Learning Experiences and Tools for AI Literacy, AI for K-12 Education, and AI, Ethics, and Public Policy.

During the panels and the affinity mapping activity (see Workshop Structure), we will engage participants in ideation and discussion surrounding questions such as: *What is AI literacy? How do we measure AI literacy? How can we ensure that we broaden AI literacy in equitable, inclusive, and culturally responsive ways? What techniques or approaches have been successful in broadening AI literacy? What are some key challenges to achieving broad public AI literacy? What are some key opportunities or spaces for introducing AI literacy learning interventions? What role will public policy play in shaping AI literacy? How can curricula, interactive learning experiences, and XAI systems all contribute to the goal of broadening AI literacy? What are some key aspects of AI ethics that need to be included in a definition of AI literacy?*

1.3 Workshop Goals

The goal of this workshop is to bring together researchers in the community working on a diverse range of topics that relate to AI literacy. We hope to involve an interdisciplinary group of researchers who may not otherwise have the opportunity to engage in an extended dialogue with each other about AI literacy. The workshop activities will be designed to facilitate sharing of research ideas across disciplines; engage an interdisciplinary group of researchers in discussion about the meaning of the term *AI literacy*, the competencies associated with it, and how to design to foster AI literacy; explore how AI literacy needs might differ across different user groups and contexts; interrogate how to foster equitable and culturally responsive AI literacy initiatives; and identify some of the

key challenges to broadening public AI literacy. Key goals of the workshop include:

- Providing an opportunity for participants to network and identify common research interests related to AI literacy
- Curating a body of work related to AI literacy that spans work on explainable AI, public policy, education, and design
- Mapping common themes, questions, and emerging research areas related to AI literacy in the CHI community

2 ORGANIZERS & DISCUSSANTS

The primary organizers for this workshop include Duri Long (Northwestern University), Jessica Roberts (Georgia Tech), and Brian Magerko (Georgia Tech). We have additionally recruited discussants to lead panels related to their expertise during the workshop and foster conversation between workshop participants. We have engaged discussants to ensure we have a range of expertise represented in a leadership capacity at the workshop, especially since it is intended to bring together folks working on different research topics within CHI. Discussants' primary role will be to prepare questions based on the papers in their panel and their own expertise to lead the panel and foster discussion between the panel participants and the audience. Discussants may also participate in the juried paper selection for their panel. Discussants were selected based on their relevant expertise for the panels they are leading. The discussants include Ken Holstein (Carnegie Mellon University), Upol Ehsan (Georgia Tech), Daniella DiPaola (MIT), and Fred Martin (University of Massachusetts Lowell). Biographies and roles for each of the workshop organizers and discussants are included below.

2.1 Duri Long, Assistant Professor, Communication Studies, Northwestern University

Dr. Duri Long is an Assistant Professor in the Department of Communication Studies at Northwestern University, where she runs the Creative Interfaces Research + Design Studio. She is a human-centered AI researcher interested in issues surrounding AI literacy and human-AI interaction. Duri's research looks at how humans interact and learn as a way of informing the design of public AI literacy interventions as well as the development of AI that can interact naturally and improvise creatively with people in complex social environments. She employs a variety of methodologies and theoretical frameworks in her research, drawing on the learning sciences, design research, and cognitive science. She has experience working with artists and museums around the country to develop co-creative, embodied exhibits and art installations involving AI and technology. Duri holds a PhD in Human Centered Computing from Georgia Tech and degrees in Computer Science and Dramatic Art from the University of North Carolina at Chapel Hill.

Dr. Long will be the primary point of contact for workshop organization and will serve as one of the primary organizers in charge of running the website, disseminating the call for papers, selecting papers, communicating with participants, and planning the workshop structure and activities.

2.2 Jessica Roberts, Assistant Professor, School of Interactive Computing, Georgia Tech

Dr. Jessica Roberts is an assistant professor in the School of Interactive Computing where she runs the Technology-Integrated Learning Environments (TILEs) Lab. She holds a PhD in Learning Sciences from the University of Illinois-Chicago with a concentration in geospatial analysis and visualization and a B.S. from Northwestern University with a concentration in theatre design. Her research examines how people learn through, with, and about data in out-of-school environments such as museums and citizen science activities and how interactive technologies mediate social, informal learning experiences. Her work on the design of interactive learning technologies has been exhibited at venues including the Field Museum of Natural History in Chicago and the New York Hall of Science. Prior to joining Georgia Tech, Dr. Roberts conducted postdoctoral research at the Human-Computer Interaction Institute at Carnegie Mellon University and the Tidal Lab at Northwestern University.

Dr. Roberts will serve as one of the primary organizers in charge of running the website, disseminating the call for papers, selecting papers, communicating with participants, and planning the workshop structure and activities.

2.3 Brian Magerko, Professor, Digital Media, Georgia Tech

Dr. Brian Magerko is a Professor of Digital Media, Director of Graduate Studies in Digital Media, & head of the Expressive Machinery Lab at Georgia Tech. He received his B.S. in Cognitive Science from Carnegie Mellon (1999) and his MS and Ph.D. in Computer Science and Engineering from the University of Michigan (2001, 2006). His research explores how studying human and machine cognition can inform the creation of new human/computer creative experiences. Dr. Magerko has been research lead on over \$18 million of federally funded research; has authored over 100 peer reviewed articles related to computational media, cognition, and learning; has had his work shown at galleries and museums internationally; and co-founded a music-based learning environment for computer science—called EarSketch—that has been used by over 1 million learners worldwide. Dr. Magerko and his lab's work have been shown in the New Yorker, Washington Post, the Smithsonian Museum of American History, USA Today, CNN, Yahoo! Finance, NPR, and other global and regional outlets.

Dr. Magerko will serve as one of the primary organizers in charge of running the website, disseminating the call for papers, selecting papers, communicating with participants, and planning the workshop structure and activities.

2.4 Ken Holstein, Assistant Professor, Human-Computer Interaction, Carnegie Mellon University

Dr. Ken Holstein is an Assistant Professor of Human-Computer Interaction at Carnegie Mellon University. In his research, he studies how AI systems are being designed and used to augment worker practices today, across a range of real-world contexts. Drawing upon approaches from HCI, AI, design, statistics, and the learning

sciences, he explores the development of new methods and tools to support more responsible AI design and deployment in practice.

Dr. Holstein will serve as a discussant for the panel on *Designing AI Literacy Learning Experiences & Tools*.

2.5 Upol Ehsan, Ph.D. Candidate, School of Interactive Computing, Georgia Tech

Upol Ehsan is a doctoral candidate in the School of Interactive Computing at Georgia Tech. Existing at the intersection of AI and HCI, his work focuses on explainability of AI systems, especially for non-AI experts, and emerging AI Ethics issues in the Global South. He is also an affiliate at the Data & Society Research Institute. His work received multiple awards at ACM CHI and HCII. His work has pioneered the notion of Rationale Generation in XAI and also charted the vision for Human-centered XAI. Along with serving in multiple program committees in HCI and AI conferences (e.g., DIS, IUI, NeurIPS), he was the lead organizer for the 2021 and 2022 CHI workshops on Human-centered XAI.

Upol Ehsan will serve as a discussant for the panel on *Explainable AI*.

2.6 Daniella DiPaola, Ph.D. student, MIT Media Lab

Daniella DiPaola is a Ph.D. student in the Personal Robots Group at the MIT Media Lab, where her research intersects artificial intelligence, education, and public policy. She conducts long-term interaction studies to understand children's perceptions of AI and designs tools that aim to give K-12 students more knowledge and agency in their use of AI technologies, including curricula addressing the social and legal implications of AI.

Daniella DiPaola will serve as a discussant for the panel on *AI, Ethics, and Public Policy*.

2.7 Fred Martin, Professor, Computer Science, University of Massachusetts Lowell

Dr. Fred Martin is professor of computer science and associate dean for teaching, learning, and undergraduate studies for the Kennedy College of Sciences at the University of Massachusetts Lowell. Martin's research group, the Engaging Computing Group, develops and studies novel computational design environments for learners, empowering them to create meaningful, personally satisfying projects. Martin was a recipient of the 2022 AAAI/EAAI Outstanding Educator Award for his contributions to the Artificial Intelligence (AI) for K-12 Initiative (aik12.org). Martin is a past chair of the Computer Science Teachers Association (CSTA) and served on the panel which produced the 2016 Digital Literacy and Computer Science (DLCS) K-12 standards for Massachusetts.

Dr. Martin will serve as a discussant for the panel on *AI in K-12 Education*.

3 WEBSITE

We will host a website explaining the purpose of the workshop, advertising a call for submissions, and posting the workshop schedule. We will also use the website as a venue for public dissemination of the workshop outputs, including a repository of the accepted

papers, a summary report of the panel discussions, and a concept map illustrating the output from the affinity mapping activity.

We will additionally use the website to link to relevant resources and demos/tools related to AI literacy, such as the AI4K12 repository and any demos/curricula shared by participants. This will create a shared repository of resources that can be used by the broader community. The URL is: <https://sites.gatech.edu/chi2023ailiteracy/>

4 PRE-WORKSHOP PLANS

We are inviting submissions on a variety of topics related to AI literacy (see Topics of Discussion). We are accepting submissions in any of the following formats:

- 2-4 page position papers
- 5-7 minute video presentations or demos with subtitles

Submissions are intended to help the organizers to group participants into panels, provide fodder for discussion amongst participants, ground the panel discussions, provide content for the discussants to pull on to lead their panel discussions, and serve as one of the workshop outputs. They will *not* be presented in a traditional paper presentation format.

If the workshop is accepted, we will post a call for papers shortly after notification. The call for papers will be posted on the website and distributed via our network and appropriate related channels and communities both within and beyond the typical CHI community (e.g., AI4K12, AAAI, learning sciences). Submissions will be due by February 13 and participants will be notified of acceptance by early March. A maximum of five papers will be accepted per panel, for a total maximum of 25 panelists (in addition to the organizers and discussant). In case co-authors of a paper wish to attend, we will allow up to 30 participants in total.

Papers will be selected by a jury consisting of the workshop organizers, discussants who choose to be involved in curating their panel, and external jury members with relevant expertise recruited on an as-needed basis. We are pursuing a juried selection approach as opposed to a blind peer review because the goals of the workshop are to curate a conversation surrounding AI literacy and we want to ensure that a variety of perspectives and research areas are represented.

We will share accepted papers and names of workshop participants and discussants on the website in advance of the workshop to enable participants to look over and potentially discuss them in advance of the workshop. Posting the papers online will additionally ensure the papers are asynchronously accessible for individuals who cannot attend the workshop in person.

5 WORKSHOP MODALITY & ASYNCHRONOUS ENGAGEMENT

The workshop will be held in person. We believe that the in-person modality will be more conducive to fostering lively interdisciplinary discussion and building community between workshop participants. In addition, we feel the planned affinity mapping activity works best when all participants are engaging via the same modality.

We will provide multiple avenues of asynchronous engagement for individuals interested in the workshop content that cannot attend synchronously in person. We plan to record video of the panel discussions and post the videos on our website for asynchronous

viewing. We will generate a summary report based on the videos and share this report on the website after the workshop to enable asynchronous engagement with the discussions that took place in the workshop. We will also share an accessible version of the concept map output from the affinity mapping activity publicly on the website. In addition, the wider community will be able to engage with the submissions before, during, and after the workshop.

6 WORKSHOP STRUCTURE

The workshop will be structured as a series of panels followed by an affinity mapping activity. This section provides more detail on the panel structure and affinity mapping activity.

6.1 Panel Discussions

As discussed in Pre-Workshop Plans, we will be accepting paper submissions to the workshop. Authors of accepted submissions will be invited to be on a panel related to their paper topic. There will be a maximum of five papers included in each panel (less if we receive fewer submissions). Each panel will be led by a discussant with relevant expertise (see Organizers & Discussants), who will facilitate a discussion amongst panelists and other workshop participants. Panels will *not* be a series of traditional paper presentations—rather, they will be a facilitated discussion designed to elicit conversation surrounding the workshop goals (see Workshop Goals). The submitted papers will serve to ground the conversation and provide fodder for discussion and will also be shared online as one of the workshop outputs.

Panel topics include:

- *Explainable AI*: This panel will focus on explainable AI (XAI) systems that can foster learning-through-interaction, particularly with non-expert human partners. Discussion in this panel will explore how XAI can contribute to broadening public understanding of AI. This panel will be led by Upol Ehsan.
- *AI in K-12 Education*: This panel focus on AI learning interventions in K-12 educational environments (e.g., curricula, activities, professional development, and tools for K-12 classrooms). This panel will be led by Dr. Fred Martin.
- *Designing Learning Experiences & Tools for AI Literacy*: This panel will focus on research related to designing interactive tools and learning experiences to support AI education and AI literacy, including public installations, informal learning experiences, and AI literacy interventions for adult users. This panel will be led by Ken Holstein.
- *AI, Ethics, and Public Policy*: This panel will focus on fostering discussion surrounding AI and ethics, particularly as this relates to public policy (e.g., developing policies to promote public understanding of AI, exploring AI literacy needed to shape public policy, creating learning experiences to help people understand ethics and policy). This panel will be led by Daniella DiPaola.

6.2 Affinity Mapping

The panel discussions will be followed by a two-hour affinity mapping activity at the end of the day. This activity is intended to aid in synthesizing ideas raised during the panel discussion and identify

common threads amongst the many research areas represented by this workshop. The output of this activity will be an affinity diagram that maps emerging questions, common themes, and future research directions related to AI literacy in the CHI community.

Affinity mapping is an activity that is commonly used in user research and analysis to sort qualitative information into groups that reflect common themes. During the panel discussions, we will ask participants to jot down notable comments, ideas, or questions from the discussions that are related to the workshop goals and topics of discussion on post-it notes. Discussants and organizers may also generate post-it notes. We will give participants additional time after the panel discussions are complete to generate notes. Participants will then be asked to work in small groups to group their post it notes based on common themes. We will then have the small groups come together and will work as a larger group to merge the themes into one summative affinity diagram on the wall of the room (or a set of easels).

The final diagram will summarize emerging themes of research and common questions, techniques, challenges, and successes across AI literacy research. This diagram will be photographed, documented, and made into an accessible format that will be shared on the website.

6.3 Proposed Schedule & Material Needs

The proposed schedule for the in-person workshop is as follows. This is subject to change (in particular, timing of coffee breaks and lunch may be shifted to match the conference schedule).

- 8:30-9AM: Intro to Workshop
- 9-10AM: Explainable AI Panel
- 10-10:30AM: Coffee Break
- 10:30-11:30AM: AI in K-12 Education Panel
- 11:30AM-12:30PM: Lunch
- 12:30-1:30PM: Designing Learning Experiences & Tools for AI Literacy Panel
- 1:30-2:00PM: Coffee break
- 2:00-3:00PM: AI, Ethics, and Public Policy Panel
- 3:00-5:00PM Affinity mapping activity

Materials needed for the workshop include post-it notes, markers, and a surface that post-it notes can be adhered to (e.g., a wall or several easels). In addition, we would appreciate AV equipment to video record the panel discussions and a projector. There will not be traditional paper presentations to project, but we may project images or project demos during the workshop if relevant to the panel discussions. We will need to be in a room large enough and with seating for 40 people (this number may be smaller if we receive fewer submissions), ideally with a setup with chairs at the front where panelists and discussants can sit.

7 ACCESSIBILITY

We will ensure that submissions are accessible in accordance with the CHI paper and video accessibility guidelines. In addition, we will ensure that all workshop outputs (e.g., videos of panel discussions, summary report) are accessible and include transcription/subtitles and figure descriptions. This will include generating an accessible version of the concept map that results from the affinity mapping

activity. Finally, the options we provide for asynchronous engagement should ensure that the workshop content is accessible to individuals who are unable to engage synchronously in person.

8 POST-WORKSHOP PLANS

After the workshop, we will post the workshop outputs to the website, including the concept map from the affinity mapping activity, a summary report of the discussions, and videos of the panel discussions. We will disseminate the workshop outputs via relevant channels in HCI, learning sciences, and AI communities.

We will gauge whether there is interest amongst the workshop participants in establishing a communication channel (e.g., mailing list or Slack channel) for communicating and sharing work after the workshop. If there is significant interest, we will set up such a channel for participants (as well as other interested parties in the CHI community) to continue to engage in discussion. If there is not interest, at the very least we will (with permission) share participants' contact information with each other.

9 CALL FOR PARTICIPATION

The following call for participation will be shared on our workshop website. We will share the call via relevant channels in HCI, learning sciences, and AI communities.

Fostering public AI literacy (i.e. “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”) has been a growing area of interest at CHI for several years, and a substantial community is forming around issues such as teaching children how to build and program AI systems, designing learning experiences to broaden public understanding of AI, developing explainable AI systems, understanding how novices make sense of AI, and exploring the relationship between public policy, ethics, and AI literacy.

This workshop aims to bring together researchers in the community working on a diverse range of topics that relate to AI literacy. We hope to involve an interdisciplinary group of researchers who may not otherwise have the opportunity to engage in an extended dialogue with each other about AI literacy. The workshop activities will be designed to facilitate sharing of research ideas across disciplines; engage an interdisciplinary group of researchers in discussion about the meaning of the term AI literacy, the competencies associated with it, and how to design to foster AI literacy; explore how AI literacy needs might differ across different user groups and contexts; interrogate how to foster equitable and culturally responsive AI literacy initiatives; and identify some of the key challenges to broadening public AI literacy.

We welcome submissions related to the following topics:

- designing formal or informal learning experiences to foster learning about AI (for adults and/or children)
- developing instruments to evaluate learner understanding of AI
- designing tools or curricula to teach about AI
- developing explainable AI (XAI) systems that can foster learning-through-interaction with non-expert human partners

- understanding how non-experts make sense of AI systems and tools
- investigating how novices learn about AI through activities such as information foraging
- designing public art or installations to engage people in learning about AI
- community-based research investigating relevant issues, concerns, and needs related to AI literacy
- developing equitable and culturally responsive AI literacy initiatives
- shaping or studying public policy and ethics related to AI education and AI literacy

More details about the workshop are available on our workshop website: <https://sites.gatech.edu/chi2023ailiteracy/> Please submit by February 13, 2023 in order to be considered for participation in the workshop. Submissions should be in one of the following formats. All submissions should comply with CHI accessibility guidelines for papers and videos. Submissions should be submitted through the PCS site for this workshop.

- 2-4 page position papers written in the new single-column ACM Master Article Template (please refer to CHI 2023 Publication Formats guide), including references
- 5-7 minute video presentations or demos with subtitles

We will use a juried selection process to decide which submissions will be accepted, meaning that a jury of experts will review the

submissions and select a subset of the submissions that will foster varied discussion and bring a variety of perspectives to the workshop event.

At least one author of each accepted submission must attend the workshop **in person** and register for the workshop and for at least one day of the conference. Workshop participants will be grouped into panels and a discussant will lead panelists in a discussion of their work and AI literacy. Submissions are intended to help the workshop organizers to group participants into panels, provide fodder for discussion amongst participants, ground the panel discussions, provide content for the discussants to pull on to lead their panel discussions, and serve as one of the workshop outputs. They will *not* be presented in a traditional paper presentation format.

Video recordings of workshop discussions, a summary report, and workshop submissions will be made available publicly on the workshop website for individuals in the CHI community who are interested in the workshop content but are unable to attend in person.

REFERENCES

- [1] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 ACM Conference on Human Factors in Computing Systems (CHI 2020)*, ACM.
- [2] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What should every child know about AI? *Proceedings of the 2019 Conference on Artificial Intelligence*, Association for the Advancement of Artificial Intelligence.