

Beyond AI Hype: A Hands-on Workshop Series for Enhancing AI Literacy in Middle and High School Students

Chinasa T. Okolo

Center for Technology Innovation
The Brookings Institution
Washington, D.C., United States
cokolo@brookings.edu

ABSTRACT

The increasing usage of AI in high-stakes decision-making underscores a pressing need for various stakeholders to understand AI, learn how to identify AI-generated content, and become aware of its societal risks. We detail outcomes from engaging underrepresented secondary school students in a 5-day workshop series consisting of brief lectures, hands-on activities, and short research assignments. We find that the workshop improved students' knowledge about AI and the ethical implications of using these technologies. Our work highlights policy implications and outlines actionable efforts needed to advance AI literacy, with the workshop content being developed into an open-source AI literacy curriculum.

CCS CONCEPTS

• **Applied computing** → **Education**; • **Computing methodologies** → **Artificial intelligence**; *Machine learning*; • **Human-centered computing**; • **Social and professional topics** → **Computer science education**; **Computing literacy**;

KEYWORDS

AI literacy; computing education; equity; human-centered AI; generative AI; technology ethics; AI pedagogy

ACM Reference Format:

Chinasa T. Okolo. 2024. Beyond AI Hype: A Hands-on Workshop Series for Enhancing AI Literacy in Middle and High School Students. In *Proceedings of the 2024 RESPECT Annual Conference (RESPECT 2024)*, May 16–17, 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3653666.3656075>

1 INTRODUCTION

For the first time in recent history, artificial intelligence (AI) has become a public conversation. From ChatGPT, which sparked the “AI revolution,” to image generators like Stable Diffusion and DALL-E, people worldwide have begun actively engaging with AI. Despite this, many people are still passive consumers and sometimes unwilling targets of these systems. Additionally, despite researchers highlighting issues of bias in AI, many of its risks are still not

understood [19]. This lack of understanding could exacerbate discrimination and misinformation from AI systems, harming users. As AI development is primarily centralized in the hands of powerful tech companies in the West [36], it will also be essential to understand how to engage marginalized populations to become equitable stakeholders in AI. Such efforts could empower these users to understand the impacts of AI, regain agency from automated decision-making, and potentially seek recourse from harms committed by these technologies.

Given the ability of AI to make essential decisions in fields such as healthcare, agriculture, and education, there is a pressing need for various stakeholders to understand AI, learn how to identify AI-generated content, and become aware of its risks to society. Improving AI literacy is just one way to do this. Researchers have defined AI literacy as “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” [25]. While there exist no prominent AI literacy frameworks to aid in facilitating pedagogy, there is an emerging area of work focused on developing AI literacy in primary and secondary school educational contexts [24, 28, 32, 47, 50].

To ensure that various stakeholders understand the benefits, limitations, and ethical implications of AI, much investment is needed to support AI literacy efforts. Creating the social and technical infrastructure to support AI literacy will serve public interests and shift the democratization of AI to include novice technology users. To progress toward this goal, the research team conducted a 5-day workshop series with underrepresented secondary school students in a rural county in New York State to engage them in various aspects of AI, ranging from data collection to algorithmic development to the ethical considerations of generative AI. We found that incorporating critical reflection assignments along with short lectures, videos, and experimentation with AI tools allowed students to understand their role as consumers of AI and how they can thoughtfully engage with these systems.

The content from these workshop sessions was adapted into an open-source curriculum, “Beyond AI Hype”, which is available online¹. The lesson plans and workshop slides are located in the accompanying GitHub repository.

2 DESCRIPTION OF PRACTICE

To answer the research question that guided the goals of this workshop series, we conducted five (5) workshop sessions across two weeks in July 2023. This section details our participants, workshop curricula, and general findings from engaging students.

¹<https://github.com/chinasatokolo/BeyondAIHype>

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 the author(s) 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.

RESPECT 2024, May 16–17, 2024, Atlanta, GA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0626-4/24/05

<https://doi.org/10.1145/3653666.3656075>

2.1 Participants

To pilot this work, the workshop team recruited a group of local middle and high school students through the Science and Technology Entry Program-Upward Partnership (STEP-UP) Summer Internship program directed by graduate students at Cornell University. The program aims to positively impact the middle/high school student population in the university area through outcome-based programming, scientific exposure, diversity initiatives, and interactions with college and graduate students. Each of the four workshop participants identified as female. One of the participants identified as Black, one as mixed-race, and two as Asian. One of the students was entering the final year of middle school (8th grade), one was entering high school (9th grade), and the other two had completed one year of high school.

2.2 Workshop Goals

To help guide the implementation of the workshop series, we defined several goals:

- Students participating in this 5-day workshop series will learn about various aspects of AI, including the socio-ethical implications of emerging AI methods such as generative AI.
- Students will critically engage with these methods and better understand how to use AI responsibly.
- Students will also improve their AI literacy by learning to recognize AI-generated content and understanding the shortcomings of AI tools like large language models and text-to-image models.

2.3 Workshops

Throughout five 3-hour sessions, workshop participants engaged in hands-on activities, brief lectures, and short research assignments. Below, we detail the aims of the workshops and the content provided in each session. The primary author has developed formal lesson plans and lecture slides from the workshops, which have been open-sourced through GitHub.

2.3.1 Session I: General Introduction to AI. The first session of this workshop series provided students with a general introduction to AI. We began the session by engaging students in a group discussion, where they shared their knowledge of AI and aspects they were interested in learning more about. To guide this discussion, students were asked three questions:

- What do you know about AI?
- What do you want to know about AI?
- What do you think is important for students your age to know about AI?

From these discussion questions, we found that students had basic knowledge about AI, stating that the acronym stood for “artificial intelligence” and that AI is embedded in robots and computers. They were interested in learning about the importance of AI, its ethical implications, how it functions, and how to use these technologies appropriately. Our workshop participants also thought it was important for students their age to know that AI has many exciting features, but it is still a work in progress. They also felt that it was important for their peers to be able to distinguish between human-generated and AI-generated content and understand

if AI could “take over the world.” The responses to the questions posed during the group discussion were used to shape the content introduced in subsequent workshop sessions.

Then, students were introduced to the AI concepts through a short lecture. First, we provided definitions of artificial intelligence and machine learning, illustrating the relationship between these two fields. We also provided examples of other subfields of AI, like deep learning, computer vision, and natural language processing. We then detailed the history of AI, starting from work done in the 1950s by Alan Turing and leading to present-day advancements in generative models and smart assistants. We also provided example use cases of AI, so students could understand how these methods are applied in real life. We also briefly introduced students to the implications of AI, which was then followed up more thoroughly in Session III.

Examples of AI

- Recommendation systems
- Self-driving cars
- Voice recognition
- Autofill and predictive text
- Face recognition



Figure 1: A lecture slide from Workshop Session I detailing example applications of AI in real-world contexts.

fig:session1

AI Activity To give students first-hand experience with AI, the next section of the first workshop incorporated the “Discover AI in Daily Life” lesson from Google [44]. This activity involved (1) students choosing three words and translating them using Google Translate into a language of their choosing, (2) using Google Auto-Draw to create an AI-generated image of each word, and (3) creating a Google Slides presentation with these words and AI-generated images. After presenting their slides individually, we debriefed on the activity and transitioned into a brainstorming session where students discussed the following questions:

- How do you use AI in your daily life?
- What are some interesting things you could do with AI?
- What are some potential use cases of AI?

2.3.2 Session II: Data and Algorithms. In the second workshop session, students learned about the importance of data and algorithms in machine learning (ML). We began the session with a group discussion on their knowledge of these concepts. To guide the discussion, students were asked the following questions:

- What do you know about datasets?
- What do you know about algorithms?
- Do you know how they impact ML models?

Then, students were introduced to the concepts of data, algorithms, and model training through a brief lecture. First, we defined data

and datasets, providing examples of popular datasets used in machine learning. We then defined algorithms and provided examples of what algorithms can be used for. To lead up to the full lecture on algorithmic bias, which would take place in Session III, students were introduced to topics of dataset and algorithmic bias and watched a short video explaining how bias impacts ML models [16]. We then explained how data is used in machine learning, reviewing the collection, cleaning, and labeling processes. We ended the lecture briefly, explaining how models are trained and evaluated.

Algorithm Examples

Algorithms can be used to train computers and robots to perform different tasks such as recognizing images, performing specific tasks, and analyzing trends from data.



Figure 2: A lecture slide from Workshop Session II detailing examples and functions of algorithms.

fig:session2

ML Activity. Next, students participated in a hands-on activity to create a simple ML model using the “Teachable Machine” platform developed by Google [17]. In this activity, students gathered data of their choosing (images, sounds, poses, etc.), trained a machine learning model, exported it, and wrote a short response to their experimentation. We directed students to test what data their models work on, test what data their models don’t work well on, strategies they used to “trick” their trained models, and how tweaking the epochs, batch size, and learning rate improves model performance. After students presented their reports, we ended the session by debriefing on the activity.

2.3.3 Session III: Algorithmic Bias. In the third workshop session, students learned about the ethical implications of AI. Similar to the previous sessions, we began the session with a group discussion on their existing knowledge of the session topic, in this case, covering bias and its impact on ML models. To guide the discussion, students were asked the following questions:

- What do you know about bias?
- What are some examples of bias?
- Do you know how bias impacts ML models?

We then presented information about bias and case studies regarding the ethical implications of AI in a brief lecture. First, we defined bias and presented various types of bias, including confirmation, stereotyping, availability, implicit, anchoring, and in-group bias. Next, we specifically presented topics relating to bias in machine learning, covering dataset bias and algorithmic bias. We also incorporated two separate videos to help students further visualize these concepts [14, 16].

Algorithmic Bias

- Algorithmic bias refers to the presence of unfair or discriminatory outcomes in automated decision-making systems or algorithms
- Biased datasets can also lead to biased algorithms




Figure 3: A lecture slide from Workshop Session III detailing algorithmic bias.

fig:session3

Bias Case Study. To lead into the activity for this workshop session, students watched a video on algorithmic bias that detailed how bias shows up in search results [3]. In this activity, students were instructed to find examples of AI bias incidents covered in the media or in research articles and create a short write-up on one example. To guide their writing process, we asked students to detail *what* their bias case study is about and what field it impacts (medicine, finance, hiring, etc.), *who* (company, researchers, etc.) developed the biased tool or algorithm, *when* the bias incident occurred, *how* it affected users, and *why* they think it is an important issue. After students presented their case studies, we moved on to a short lecture about the ethical issues of data annotation work, presenting a video on data labeling [43] and asking students to read a recent article on workers involved in labeling data for ChatGPT [39]. We then ended the session by debriefing on the activity.

2.3.4 Session IV: Introduction to Generative AI. In the fourth workshop session, students learned about the basics of Generative AI. To begin this session, we initiated a group discussion on the students’ knowledge of generative AI. To guide the discussion, students were asked the following questions:

- What do you know about generative AI?
- What are some examples of generative AI?

Then, students were introduced to the concepts of generative AI in a lecture. First, we defined generative AI and shared a timeline of developments in generative models from the 1950s to the present. Next, we shared examples of how generative AI is being used (image generation, question answering, music composition, etc.) and transitioned to the next topic of large language models (LLMs) and text-to-image models. After defining LLMs, the instructor detailed how they function and shared examples of popular models like ChatGPT, Bard, Midjourney, and Dall-E. To codify these concepts, students watched a short video from The Economist [13]. The last section of the lecture discussed the limitations of generative AI and examples of gender, cultural, and regional bias produced by these systems. To transition into the activity for this session, students read an article on gender bias from AI text-to-image generators [33] and watched a video from Bloomberg detailing a recent study on racial and gender bias in generative AI [45].

Functions of LLMs

LLMs can generate relevant responses given a prompt or input. They can engage in conversations, answer questions, provide explanations, and perform a variety of language-related tasks.




Figure 4: A lecture slide from Workshop Session IV detailing functions of large language models (LLMs).

fig:session4

Generative AI Activity. To give students more context into how generative AI models work, the first part of the activity for this session involved experimentation with a tool developed by the New York Times to detail the process of training an LLM “from scratch” [5]. Students could select literature from Jane Austen, Shakespeare, Moby-Dick, Star Trek, Harry Potter, and the Federalist Papers and go through an interactive demonstration to learn how LLMs “learn” language from 250 to 30,000 rounds of training on these texts. Next, to help students develop strategies to discern between original and AI-generated content, students engaged with an interactive quiz developed by the New York Times to test if they could differentiate if specific text samples were written by a fourth-grade student or AI-generated. For the last part of this activity, students played an interactive game to understand how LLMs can be manipulated to reveal information through a technique known as “prompt injection”, where users manipulate LLMs to perform unintended actions involuntarily [42]. For each section of this activity, students were asked to write up their experiences interacting with these tools and gave short presentations. We then ended the session by debriefing on the activity.

2.3.5 Session V: Implications of Generative AI. In the fifth and final workshop session, students reviewed concepts about Generative AI and experimented with popular AI tools to generate images and text. To begin this session, we initiated a group discussion on the students’ experiences using generative AI tools. To guide the discussion, students were asked the following questions:

- What do you know about ChatGPT or other generative AI tools?
- Have you experimented with generative AI tools?
- How was your experience with these tools?

Next, we reviewed the concept of generative AI by watching the first ten minutes of a video developed by Google [8] and transitioned into the planned activity.

Ethical Implications of Generative AI Activity. To give students an opportunity to increase their direct engagement with

generative AI tools, students experimented with ChatGPT and Microsoft Image Generator. In this experimentation, students generated text and image content and wrote up a short reflection analyzing what these tools do well and the kinds of tasks they appear to perform poorly at. Students then presented their findings. Next, the students watched a video on the ethical implications of ChatGPT [31] and read an article from the New York Times on university responses to AI chatbots [21]. After engaging with this content, students were asked to create a short write-up guided by the following questions:

- Why do you think many educators are worried about ChatGPT and similar technologies?
- Who is responsible for ensuring these tools are used responsibly?
- When should Generative AI not be used?
- How do you think students can misuse ChatGPT and other technologies?
- Why is it important to use generative AI responsibly?

To gain further perspective on the ethical concerns of using generative AI tools in the classroom, students watched a video detailing how a school in Hong Kong embraced this technology [41]. Finally, to incorporate what they learned from the workshop sessions, we engaged students in a discussion about the implications of using generative AI and how these methods could be ethically used to complement their learning. We then ended the last session by asking students to fill out a post-workshop survey, debriefing on the workshop, and providing opportunities for the students to ask any lingering questions.

2.4 Analysis

Pre-Workshop Survey. To guide the development of the workshop materials, students were asked to fill out a pre-workshop survey before the start of the first session. This survey asked the students what they knew about AI, if they had used any AI tools, their understanding of AI (using a Likert scale rating from 1-5, with 1 indicating low AI understanding and 5 indicating high AI understanding), their familiarity with the ethical implications of AI, and what they were specifically interested in learning about AI.

When asked about their AI knowledge, all of the students generally knew that “AI” stood for “artificial intelligence,” with one student, in particular, mentioning that it aims to “*imitate human knowledge and can generate an answer to almost anything you ask it.*” Students often mentioned their knowledge of AI in relation to text-generation applications like ChatGPT. When rating their understanding of how AI works, the students had an average rating of 3 with a median rating of 3.5. Exactly half of the students responded that they were familiar with the ethical implications of AI. When asked what they were specifically interested in learning about AI, students mentioned wanting to understand how AI functions, examples of real-world applications of AI, the origin of AI, and the potential implications of this technology.

Post-Workshop Survey. To understand the impact of the workshop on their perceptions of AI, students also filled out a post-workshop survey. This survey asked them what they learned about AI, their interest in using AI tools in the future based on their workshop experiences, how the workshop impacted their understanding

of AI and its associated ethical implications, and any remaining questions they had about AI.

When asked about their takeaways from the workshop, students mentioned that they learned how AI is developed, what it can be used for, and its ethical implications. Many of the students emphasized their improved perceptions of AI, especially as it pertains to the limited ability of AI to become sentient or “take over the world.” Students also recognized that AI is increasingly being used in various aspects of their lives and that there are many implications to using AI. All of the students responded positively to using AI tools in the future. When rating if their understanding of how AI works improved (also scored using a Likert scale rating from 1-5, with one indicating “strongly disagree” and five indicating “strongly agree”), students provided an average rating of 4.75 with a median rating of 4.5.

2.5 Outcomes

We found unique insights from engaging the students in this 5-day workshop series. First, while all of the students currently engaged with AI tools such as ChatGPT, all of them had a limited understanding of AI concepts (i.e., recommendation systems, data labeling, large language models), which was not unexpected given their respective educational levels. However, before engaging in the workshops, some students were somewhat familiar with the implications of AI from exposure to personal readings or popular news topics. Our post-workshop discussions and survey revealed that the workshop participants had various interests in understanding the inner workings of AI. For example, one student mentioned wanting to know what *“the code [for AI models] actually looks like”* and another wanted to understand how to actually build AI models, aside from what they had experimented with through Teachable Machine. Our workshop also provided an opportunity to expose students to the harms of AI, especially as it pertains to the labor involved in the production and annotation of data, an essential aspect of training AI systems. From this exposure, we found that critically engaging students in interactive lectures while providing them with the space to reflect on their personal usage of AI tools improved their understanding of AI and their ability to recognize potential harms.

3 POSITIONALITY

The workshop team is based in the United States and has developed curricula for and conducted fieldwork with underserved communities in low-resource regions within the United States and the Global South. The primary author identifies as female and is from an underrepresented racial background (in the context of the United States). The author also has 8+ years of experience teaching computing and AI topics to students ranging from kindergarten to undergraduate levels. As AI researchers whose work centers on marginalized communities, we believe in elevating the voices of local populations and actively including them in AI development. We approach our research through an equity-driven [20] and emancipatory action [4, 22] mindset, where we aim to identify the opportunities and challenges of improving AI literacy in marginalized populations while underscoring the needs of those who will interact with or be subjects of AI-enabled technologies.

4 LIMITATIONS AND ASSUMPTIONS

This project is subject to several potential limitations. The primary constraint of this project lies in our small participant sample size, as the program the workshop series was a part of intentionally targeted a small number of students for each workshop. Additionally, all the students who participated in the workshop were female, so our preliminary findings may not generalize across students of different genders. We found that our students also had many intersectional identities (e.g., Black, rural, female) but could not make any correlations on the impact of participant backgrounds on AI understanding and learning outcomes due to our sample size. However, we find that this is an important area for future work, given the burgeoning amount of research examining the impact of the digital divide on educational outcomes for underrepresented students, particularly in light of the COVID-19 pandemic [7, 10, 15, 37].

Given that there are no standardized methods for measuring AI literacy, we also acknowledge that students may have overestimated their AI knowledge in the pre-workshop survey. We also found that some experimentation was needed to ensure we had enough content for the sessions because they were longer than traditional lectures (e.g., 3 hours per workshop vs. 1-hour class periods). The first workshop session was primarily lecture-based and included one activity, which wasn’t long enough for 3 hours. In subsequent sessions, we developed slightly longer lectures incorporating videos and then leveraged article readings combined with short group discussions on the assigned articles before starting the primary activity. If educators are interested in adapting these workshops to a larger group of students, it may be necessary to explore different measures to gauge students’ initial AI knowledge and track its change over time, which could help them choose what content to include for each workshop session.

Additionally, with the short lead time to produce the content for the workshops, the team found it difficult to obtain short video explainers on more complicated topics discussed in the workshop, such as generative AI, that were accessible to the age group of the participants. To get around these constraints, we often found videos geared to more technical audiences and shortened the amount of content we would show students. For example, in the last session, we showed students the first 10 minutes of a video on Generative AI [8] to briefly recap the more in-depth introduction we provided to students in the previous workshop. Since the workshops were held in June and July 2023, there has been an increase in educational content about AI topics geared towards secondary school students that may be useful to educators interested in adapting the curriculum for their respective student populations [30].

The workshop series described in this report primarily focused on teaching secondary school students, so some of the content provided in the lecture slides may not be accessible to younger students (K-5). However, considering the urgent need for AI literacy to reach a more significant number of other key stakeholders, working with teachers and students at primary and tertiary schooling levels is an essential area for future work.

5 FUTURE DIRECTIONS

The pilot of this 5-day workshop provides various avenues for future work. Future steps for this research involve scaling up the

workshops to reach more students and gain feedback (e.g., through structured surveys) on how AI curriculum could be enhanced for various age groups. While the planning and implementation of this workshop were not incorporated into a formal research study, we also plan to examine how to leverage interviews, focus groups, and co-design workshops to design relevant AI literacy curricula. Our pre- and post-workshop surveys provided preliminary insights on student learning outcomes from the workshops, but we find it necessary to make progress toward the development of standardized methods for measuring AI literacy. Understanding how to implement effective evaluation methodologies for measuring AI literacy will help provide tangible metrics that future studies can incorporate.

There are also opportunities for future versions of this workshop and the open-sourced curriculum to cater to the diverse needs of underrepresented students. We believe this curriculum can be adapted and personalized in ways relevant to students' backgrounds and interests. For example, we relied heavily on more popularized case studies of bias that primarily centered on gender and racial/ethnic bias. Still, there are many other aspects of social identity that weren't covered in depth during the sessions, such as disability or socioeconomic class. While the current workshop structure allowed for flexibility in the types of activities that students could engage in (e.g., the Teachable Machine activity provided options for students to train an image, sound, or pose classification model), there are opportunities for increased personalization. However, future work will be needed to understand how the existing workshop material can meet the unique learning needs of students and redesign appropriate curriculum for their needs.

We also find it essential to gain perspectives from educators, and future avenues for this work could involve co-designing with K-12 educators to develop AI curricula for students across primary and secondary educational settings. To help this process, immediate future steps for this work will involve examining contemporary measures that quantify AI knowledge and reviewing existing work in HCI and computer science education to inform the development of a human-centered AI framework for building AI literacy. This AI literacy framework would then guide AI curricula co-design, contributing to future efforts to advance AI literacy.

6 RECOMMENDATIONS FOR RESEARCH, PRACTICE, AND POLICY

Increasing interest in regulating AI has spurred thousands of policy proposals, white papers, and formal legislation intending to protect the rights of consumers and hold developers of AI-enabled technologies accountable [26, 35, 38]. While responsibility primarily lies on companies to comply with AI regulations, more efforts are needed to ensure that the public is aware of the increasing role of AI in everyday scenarios that range from employment screenings to social assistance disbursements to medical decision-making. Many of these systems have already caused significant harm to consumers, and fortunately, legal action against these practices has begun to become more prominent [9]. As companies and organizations increase their efforts toward integrating AI technologies into their services and products, harms against consumers will likely increase, thus making it essential for heightened efforts to protect

consumers. Along with increased regulation and oversight of industry practices, this is where AI literacy stands to make a significant impact.

While 84% of adults in the United States are digitally literate [27], the general public lacks AI literacy and has low knowledge about data privacy topics like two-factor authentication and website encryption [46]. Policy measures such as the United States Executive Order on AI emphasize protecting Americans' privacy through external means, such as developing privacy-enhancing technologies [1]. However, these methods are limited in their effectiveness unless citizens become aware of how to restrict the use of their personal data and how to prevent unintended sharing of confidential information, which is a common occurrence with AI chatbots [18, 48, 49]. Increasing data and privacy literacy is a fundamental step toward advancing efforts toward AI literacy. With this in mind, researchers should focus on expanding efforts toward understanding practices that hinder users from comprehending and enacting data privacy measures. Additionally, developers should work towards making privacy measures more accessible to users. While the enactment of GDPR has led to broader policy measures limiting website cookie tracking, many of these measures are deceptive and often lead users to "consent" to tracking [2, 6, 23, 29]. To help enforce ethical measures, governments should pass comprehensive policy measures against deceptive data privacy practices. Given that only thirteen states have passed data privacy laws [40], the United States should work towards national data privacy regulation to ensure that all citizens are protected against deceptive measures and have proper rights to privacy.

As governments move toward regulating AI, more proposals are actively advocating for AI literacy, such as the AI Literacy Act introduced in the House of Representatives in December 2023 [11]. However, there is a lack of AI curricula tailored to audiences such as K-5 students, low-literate users, and other marginalized populations. As countries race to improve AI literacy in their respective populations, much more research is needed to develop multi-faceted approaches to building AI literacy. While traditional educational settings are prime choices to begin the implementation of AI literacy efforts, there should be more accessible outlets to implement AI literacy upskilling, such as workshops at community centers, libraries, and through programs such as Job Corps in the US, which offer residential career training for young adults without a high school diploma [34]. Research also shows that existing teaching resources for AI education currently do not support the needs of educators [12]. To help increase the effectiveness of AI literacy training, further collaboration between computing education researchers and educators is needed to strengthen AI pedagogy. Efforts like the AI Pedagogy Project from the metaLAB at Harvard University have made significant progress towards such efforts by developing guides to introduce educators to AI, help them form classroom policies around AI, and providing a crowdsourced repository of assignments [30]. Researchers will also have to explore the considerations needed to effectively upskill different stakeholders (e.g., students vs. educators vs. doctors vs. machinists) who are situated in various contexts (e.g., rural vs. urban and low-resourced vs. high-resourced). Many opportunities exist to improve AI literacy

both domestically and globally; however, such efforts will need intentional collaboration from academia, industry, government, and civil society.

7 CONCLUSION

This experience report details a 5-day workshop series conducted with underrepresented middle and high school students in Upstate New York. We find that integrating critical reflection assignments along with short lectures, videos, and experimentation with AI tools into the workshop sessions provides students an opportunity to understand their role as consumers of AI and ways they can thoughtfully engage with these systems. Our work contributes to existing literature on AI literacy and computing education while also providing additional resources that educators can leverage to teach AI concepts and engage students in emerging AI topics. The increasing interest in building AI literacy in varying stakeholder populations outside of secondary school settings potentially suggests that our findings could be relevant to a broader set of populations. However, future curricula development and implementation will be required to understand how our work can generalize to other domains.

ACKNOWLEDGMENTS

The workshops were facilitated through the Cornell Science and Technology Entry Program-Upward Partnership (STEP-UP) Summer Internship program. We are thankful to the students who participated in our workshop sessions. This work was funded in part by the ICICLE-National Science Foundation AI Institutes Educational Fellowship.

REFERENCES

- [1] Biden-Harris Administration. 2023. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. *The White House* (2023). Retrieved January 15, 2024 from <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- [2] Koen Aerts. 2021. *Cookie Dialogs and Their Compliance*. Ph.D. Dissertation. Master's thesis, Open University of the Netherlands.
- [3] Algorithms and Society. 2020. Algorithmic Bias, Discover How It's Affecting You! Retrieved July 1, 2023 from https://www.youtube.com/watch?v=7SVin0_ePGQ
- [4] Jeffrey Bardzell and Shaowen Bardzell. 2015. What is Humanistic HCI? In *Humanistic HCI*. Springer, 13–32.
- [5] Aatish Bhatia. 2023. Let Us Show You How GPT Works — Using Jane Austen. *The New York Times* (2023). Retrieved July 1, 2023 from <https://www.nytimes.com/interactive/2023/04/26/upshot/gpt-from-scratch.html>
- [6] Nataliia Bielova, Laura Litvine, Anysia Nguyen, Mariam Chammat, Vincent Toubiana, and Estelle Harry. 2024. The Effect of Design Patterns on (Present and Future) Cookie Consent Decisions. In *USENIX Security Symposium*. USENIX Association. Accepted for publication.
- [7] Alyvia Bruce. 2020. Bridging the Technological Divide in Education. *Harvard Political Review* (2020).
- [8] Google Cloud. 2023. Introduction to Generative AI. Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=G2fqAlgoPo&t=1s>
- [9] Federal Trade Commission. 2023. Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology without Reasonable Safeguards. (2023). Retrieved January 15, 2024 from <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>
- [10] Christopher Cruz. 2021. From digital disparity to educational excellence: closing the opportunity and achievement gaps for low-income, Black and Latinx students. *Harv. Latinx L. Rev.* 24 (2021), 33.
- [11] Andrew Donnelly. 2023. Reps. Blunt Rochester & Bucshon Introduce Bipartisan Artificial Intelligence Literacy Bill. *Office of Representative Lisa Blunt Rochester (D-Del.)* (2023). Retrieved January 15, 2024 from <https://bluntrochester.house.gov/news/documentsingle.aspx?DocumentID=4062>
- [12] Stefania Druga, Nancy Otero, and Amy J Ko. 2022. The landscape of teaching resources for ai education. In *Proceedings of the 27th ACM Conference on Innovation and Technology in Computer Science Education Vol. 1*. 96–102.
- [13] The Economist. 2023. Generative AI: what is it good for? Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=gCDacaohqA>
- [14] Institute for Public Policy Research. 2020. A Algorithmic Bias Explained. Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=tia5OHE98F4>
- [15] Amy L Gonzales, Jessica McCrory Calarco, and Teresa Lynch. 2020. Technology problems and student achievement gaps: A validation and extension of the technology maintenance construct. *Communication research* 47, 5 (2020), 750–770.
- [16] Google. 2017. Machine Learning and Human Bias. Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=59bMh59JQDo>
- [17] Google. 2017. Teachable Machine. Retrieved July 1, 2023 from <https://teachablemachine.withgoogle.com/>
- [18] Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*. 79–90.
- [19] Alexa Hagerty and Igor Rubinov. 2019. Global AI ethics: a review of the social impacts and ethical implications of artificial intelligence. *arXiv preprint arXiv:1907.07892* (2019).
- [20] Christina Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing Community-based Collaborative Design: Towards More Equitable Participatory Design Engagements. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [21] Kalley Huang. 2023. Alarmed by A.I. Chatbots, Universities Start Revamping How They Teach. *The New York Times* (2023). Retrieved July 1, 2023 from <https://www.nytimes.com/2023/01/16/technology/chatgpt-artificial-intelligence-universities.html>
- [22] Os Keyes, Josephine Hoy, and Margaret Drouhard. 2019. Human-computer Insurrection: Notes on an Anarchist HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [23] Oksana Kulyk, Willard Rafnsson, Ida Marie Borberg, and Rene Hougaard Pedersen. 2022. "So I Sold My Soul": Effects of Dark Patterns in Cookie Notices on End-User Behavior and Perceptions. In *Proceedings of 2022 Symposium on Usable Security and Privacy*. Internet society.
- [24] Irene Lee, Safinah Ali, Helen Zhang, Daniella DiPaola, and Cynthia Breazeal. 2021. Developing middle school students' AI literacy. In *Proceedings of the 52nd ACM technical symposium on computer science education*. 191–197.
- [25] Duri Long and Brian Magerko. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.
- [26] Gianclaudio Malgieri. 2019. Automated decision-making in the EU Member States: The right to explanation and other "suitable safeguards" in the national legislations. *Computer law & security review* 35, 5 (2019), 105327.
- [27] Saida Mamedova, Emily Pawlowski, and Lisa Hudson. 2018. A Description of U.S. Adults Who Are Not Digitally Literate. *U.S. Department of Education Statistics in Brief* (2018).
- [28] Radu Marescu-Istodor and Ilkka Jormanainen. 2019. Machine learning for high school students. In *Proceedings of the 19th Koli calling international conference on computing education research*. 1–9.
- [29] Thomas Mejtøft, Erik Frängsmyr, Ulrik Söderström, and Ole Norberg. 2021. Deceptive design: cookie consent and manipulative patterns. In *34th Bled eConference-Digital support from crisis to progressive change, Online, June 27-30, 2021*. 397–408.
- [30] metaLAB at Harvard. 2024. The AI Pedagogy Project. Retrieved January 15, 2024 from <https://aipedagogy.org/>
- [31] NBC News. 2023. Breaking down ethical questions surrounding new chatbot GPT-4. Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=C02-YMqwpS8>
- [32] Davy Tsz Kit Ng, Jiahong Su, Jac Ka Lok Leung, and Samuel Kai Wah Chu. 2023. Artificial intelligence (AI) literacy education in secondary schools: a review. *Interactive Learning Environments* (2023), 1–21.
- [33] Kristina Nikolic and Jelena Jovicic. 2023. Reproducing inequality: How AI image generators show biases against women in STEM. Retrieved July 1, 2023 from <https://www.undp.org/serbia/blog/reproducing-inequality-how-ai-image-generators-show-biases-against-women-stem>
- [34] U.S. Department of Labor. 2023. Job Corps. Retrieved January 15, 2024 from <https://www.dol.gov/agencies/eta/jobcorps>
- [35] White House Office of Science and Technology Policy. 2022. Blueprint for an AI Bill of Rights. *The White House* (2022). Retrieved January 15, 2024 from <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- [36] Chinasa T. Okolo. 2023. Addressing global inequity in AI development. In *Handbook of Critical Studies of Artificial Intelligence*. Edward Elgar Publishing, 378–389.
- [37] Paul M. Ong. 2020. COVID-19 and the Digital Divide in Virtual Learning, Fall 2020. *UCLA Center for Neighborhood Knowledge* (2020).
- [38] EU Parliament. 2023. EU AI Act: first regulation on artificial intelligence. Retrieved January 15, 2024 from <https://www.europarl.europa.eu/news/en/>

- headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence
- [39] Billy Perrigo. 2023. OpenAI Used Kenyan Workers on Less Than \$2 Per Hour. *TIME* (2023). Retrieved July 1, 2023 from <https://time.com/6247678/openai-chatgpt-kenya-workers/>
- [40] F. Paul Pittman. 2023. US Data Privacy Guide. *White & Case LLP* (2023). Retrieved January 15, 2024 from <https://www.whitecase.com/insight-our-thinking/us-data-privacy-guide>
- [41] South China Morning Post. 2023. How a Hong Kong school embraces ChatGPT in the classroom. Retrieved July 1, 2023 from <https://www.youtube.com/watch?v=z3yYmuH0Ymw>
- [42] Deval Shah. 2023. The ELI5 Guide to Prompt Injection: Techniques, Prevention Methods & Tools. *Lakera AI* (2023). Retrieved August 1, 2023 from <https://www.lakera.ai/blog/guide-to-prompt-injection>
- [43] DW Shift. 2023. Doing Grueling Work for an AI: Data Labeling. Retrieved July 1, 2023 from https://www.youtube.com/watch?v=ug_p2wHhla0&t=1s
- [44] Google Applied Digital Skills. 2022. Discover AI in Daily Life. Retrieved July 1, 2023 from <https://applieddigitalskills.withgoogle.com/c/middle-and-high-school/en/discover-ai-in-daily-life/overview.html>
- [45] Bloomberg Technology. 2023. Calling Out the Racial and Gender Bias in Generative AI. Retrieved July 1, 2023 from <https://www.bloomberg.com/news/videos/2023-06-09/calling-out-the-racial-and-gender-bias-in-generative-ai-video>
- [46] Emily A. Vogel and Monica Anderson. 2019. Americans and Digital Knowledge in 2019. *Pew Research Center* (2019). Retrieved January 15, 2024 from <https://www.pewresearch.org/internet/2019/10/09/americans-and-digital-knowledge/>
- [47] Randi Williams, Stephen P Kaputso, and Cynthia Breazeal. 2021. Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 15678–15686.
- [48] Qionikai Xu, Lizhen Qu, Zeyu Gao, and Gholamreza Haffari. 2020. Personal information leakage detection in conversations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 6567–6580.
- [49] Jiahao Yu, Yuhang Wu, Dong Shu, Mingyu Jin, and Xinyu Xing. 2023. Assessing Prompt Injection Risks in 200+ Custom GPTs. *arXiv preprint arXiv:2311.11538* (2023).
- [50] Helen Zhang, Irene Lee, Safinah Ali, Daniella DiPaola, Yihong Cheng, and Cynthia Breazeal. 2023. Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. *International Journal of Artificial Intelligence in Education* 33, 2 (2023), 290–324.