

Developing Middle School Students' AI Literacy

Irene Lee
MIT STEP Lab
Cambridge, Massachusetts
ialee@mit.edu

Safinah Ali
MIT Media Lab
Cambridge, Massachusetts
safinah@media.mit.edu

Helen Zhang
Boston College
Chestnut Hill, Massachusetts
zhangzm@bc.edu

Daniella DiPaola
MIT Media Lab
Cambridge, Massachusetts
dipaola@media.mit.edu

Cynthia Breazeal
MIT Media Lab
Cambridge, Massachusetts
cynthiab@media.mit.edu

ABSTRACT

In this experience report, we describe an AI summer workshop designed to prepare middle school students to become informed citizens and critical consumers of AI technology and to develop their foundational knowledge and skills to support future endeavors as AI-empowered workers. The workshop featured the 30-hour “Developing AI Literacy” or DAILy curriculum that is grounded in literature on child development, ethics education, and career development. The participants in the workshop were students between the ages of 10 and 14; 87% were from underrepresented groups in STEM and Computing. In this paper we describe the online curriculum, its implementation during synchronous online workshop sessions in summer of 2020, and preliminary findings on student outcomes. We reflect on the successes and lessons we learned in terms of supporting students’ engagement and conceptual learning of AI, shifting attitudes toward AI, and fostering conceptions of future selves as AI-enabled workers. We conclude with discussions of the affordances and barriers to bringing AI education to students from underrepresented groups in STEM and Computing.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; *Learning paradigms; Learning settings; Machine learning approaches; Machine learning algorithms*; • **Social and professional topics** → **Computing education**; *Computational thinking; Model curricula*.

KEYWORDS

machine learning education, computing education, computational thinking

ACM Reference Format:

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 (SIGCSE '21)*, March 13–20, 2021, Virtual Event, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3408877.3432513>



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGCSE '21, March 13–20, 2021, Virtual Event, USA.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8062-1/21/03.
<https://doi.org/10.1145/3408877.3432513>

1 INTRODUCTION

With the rapid growth of enormous amounts of data and computation power, Artificial Intelligence (AI) is making unprecedented industrial and societal impact. The rapid expansion of AI across fields and industries necessitates developing a workforce with strong computational skills and specifically, the knowledge and capability to work with AI. Despite this need, recent studies, as well as our own research, have raised questions about the extent to which youth are aware of AI in their everyday lives and its application in industries of the future that may limit their interest in pursuing learning trajectories that lead toward careers in these fields.

The current interest in and call for AI activities in K-12 education echoes a more general demand for computational and scientific literacy among young people [42]. Yet AI literacy in K-12 is in its infancy and AI education tools, activities, and curricula, as well as research to ascertain their efficacy, are needed to address the demand. Several efforts are underway to build tools and curricula that enable students to interact with AI [3, 6, 14, 20, 21, 44]. The field is beginning to learn how people gain an understanding of AI concepts and processes and gain the ability to incorporate AI processes within their own applications [7, 15, 19, 39, 43, 46]. Despite its novelty, AI Literacy has been defined based on existing research [28] and National guidelines for teaching AI in K-12 are being developed by the AI4K12 Initiative [1].

In parallel to the need to build students’ AI Literacy, broadening participation in AI is of utmost importance to ensure that the design and utilization of AI technologies are inclusive and not reinforcing inequities based on demographic variables and to address historical under-representation of people of color and women in AI. Through their participation in developing the AI technologies of the future, persons from underrepresented groups in STEM and computing and their allies can work together towards ensuring that the AI industries of the future are founded in principles of inclusivity, provide equitable access, include consideration of multiple stakeholders and potential users, and minimize the potential for bias.

2 THEORETICAL FOUNDATIONS

Our approach focused on middle school-aged youth for several reasons. First, young adolescence is an important time for youth to form an identity in STEM and start thinking about their future career interests [11, 29]. Second, middle school students increasingly have daily contact with AI without realizing it. A Pew report found

that 36% of parents of a child 11 or younger report their child uses or interacts with a voice-activated assistant, such as Apple's Siri or Amazon Alexa [10]. The same report stated that by age twelve, 50% of adolescents have had exposure to AI moderated and generated content on social media websites such as TikTok, Snapchat, and Instagram. While AI-powered platforms such as TikTok and YouTube provide endless entertainment and educational opportunities, their algorithms have the potential to introduce harm [30]. It's important that students are aware of how AI technologies work in order to be responsible consumers of them.

The design of the DAlly workshop is grounded in literature on child development, ethics education, and career development. When considering middle school students' capability to learn AI concepts such as classification and prediction, we referred to literature on children's developmental trajectories. In terms of learning and cognition, middle school students are able to think more abstractly than younger students [5, 12, 22] and they can incorporate new knowledge into existing schemas [22, 23]. The mechanism by which this might occur was postulated by Piaget (2001) as "reflective abstraction" in which actions at one level become objects of reflection at the next level prior to a reorganization of concepts. Countering his earlier work suggesting that children around the age twelve become capable of learning about abstraction, in his later work Piaget wrote that children as young as 18 months can utilize abstraction and that abstraction is used continually when learning [33]. Researchers' findings on children's ability to learn abstract concepts in computer science reinforce Piaget's conjecture. [16, 40].

Another consideration was whether ethics education was appropriate for middle school students. Research has found that even young children have a conception of ethical behavior. Li et al. evidenced children under the age of 7 show a preference for self-satisfying action but around age 7, getting more while another gets less is seen as not "just" [26, 38]. Additionally, Wood found students in the middle childhood years typically have developed a sense of morality and a conscience based on values [45]. At the collegiate level, research has shown that embedding ethics throughout a course leads to an increase in students' consideration of ethics in technology design [38]. Incorporating ethical and social dimensions into science or engineering courses can increase student retention rates and better attract students, particularly underrepresented minorities and women, as they in general tend to be more drawn to topics that emphasize helping others and social concern [13, 31]. Previous studies that incorporated ethical considerations such as algorithmic bias and ethical design of recommender systems into a middle school AI and Ethics curriculum found that students were able to engage with these concepts [32].

The appropriateness of engaging middle school students in consideration of future careers is also grounded in research. Middle school years are an important time in forming youths' future career interests and pursuits in STEM [11, 29]. An exposure to STEM careers in middle school is critical because students who do not express STEM-related aspirations at age 10 are unlikely to develop them by the age of 14 [8]. Further recent theories on career development call for programs that focus on developing individuals' career adaptability at an early age [36]. Researchers argue that people craft their careers by integrating their personal needs and interests with social expectations regarding the preparation for and

participation in different job roles, and, thus, their adaptation to their environment [9, 25].

3 DEVELOPING AI LITERACY

3.1 Institutional Collaboration

The DAlly project is an interdisciplinary collaboration between Massachusetts Institute of Technology (MIT) Scheller Teacher Education Program (STEP) and Personal Robots Group (PRG) and Boston College. MIT STEP Lab contributed prototypes of participatory simulations and games that taught AI concepts such as decision trees and neural networks and expertise in teaching middle schoolers about models and modeling. MIT PRG shared a wealth of AI curricula and tools and experience designing workshops for k12 students and developing tools. Boston College's expertise in learning sciences, assessment design, and career development was critical to the research design and implementation. Our community partner provided essential support during recruitment, registration, and provided operational support during the workshop. Importantly, as a minority owned and operated organization, the organization's leaders served as a sounding board and helped to frame and facilitate discussions with youth about the negative impacts of bias in AI systems on people of color.

3.2 Target Audience

Our target audience was students of color. We recruited youth from underrepresented groups in STEM and computing through our community partner. Family information meetings were hosted online and emails were sent to interested families with middle schoolers enrolled in previous summer and afterschool programs. We reached out to parents who signed their child(ren) up for summer programs and hosted welcome meetings to inform them of the program and research study prior to seeking consent/assent. Thirty-one middle school students participated in the AI workshop. With respect to gender, 61% of participants were female, 39% were male. Of the attendees 87% were from underrepresented groups in STEM and computing fields: 58% were Black, 13% were Hispanics/Latinos, 6% were Asian females, 10% were from two or more underrepresented groups in STEM and computing, and two students (6%) did not respond. The age range was 10.4 to 14.7 years old with an average age of 12.57 years old. Students were asked about their prior experience with using Scratch as a proxy for programming/coding experience: 42% had no or little prior experience and 39% had "some" to "a lot" of experience.

3.3 Approach

Our approach is based on the following theory of action. An ethics framing, exposure to AI-enhanced careers, and the relating AI to their daily lives, current events, and real-world problems engages students from diverse backgrounds and raise their perceptions of the relevance of AI in their lives. Interacting with AI applications will whet their curiosity about how these tools work as they train and test AI models. Online participatory simulations will help students build mental models of AI mechanisms and algorithms in action and expose how bias can be embedded in AI systems through datasets. Through a series of career-focused activities, students uncover their strengths, identify jobs that both match their interests

and are aligned with their strengths, understand the occupational implications of the advancement of AI, and outline resources and barriers along their paths to their future dream jobs. Importantly, youth will recognize their AI career-related skills and resources, and develop their capacity to internalize these skills as they formulate future career ideas.

3.4 The Curriculum

The DAILY curriculum interweaves teaching of AI concepts, raising awareness of AI adoption in future jobs, and the investigation of ethical issues in AI. Learning activities include hands-on games, discussions, and building projects that integrate AI capabilities into a custom Scratch-based toolkit. Students learned key AI concepts such as Logic systems, Machine Learning, and Generative Adversarial Networks (GANS) and were engaged in working with both classification and generative uses of AI to develop a fundamental yet holistic view of AI.

The 30-hour curriculum features (1) an Introduction to AI, (2) logic systems (e.g., decision trees constructed by humans), (3) supervised learning (e.g., concepts, processes, and bias) with Google's Teachable Machine, (4) neural networks through a participatory simulation game, and (5) Generative Adversarial Networks or GANS. Within each unit students investigate the existence and causes of algorithmic bias, its societal and ethical implications and ways to mitigate bias; and also gain awareness of AI related careers, recognize their own strengths and interests for future jobs, and realize the importance of technical skills development and the ongoing nature of change and adaptation in today's job world. The detailed curriculum can be found at aieducation.mit.edu/daily [2].

Below we describe the connections made between AI concepts, ethics, and career awareness in unit 6. In this unit, we begin by introducing generative machine learning through a color generation activity called "Classifier vs. Generator". Since these students have already completed a lesson on classification using Supervised Learning, we introduce the process of generation through the creation of new media, and contrast it with the process of classification in which existing media are arranged into categories. Students experience the processes of generation and classification by mixing colors in an online platform called "TryColors" [4]. During the generation stage, students observe that even with a few input colors, they can create a varied palette of colors. Next, in the classification stage, students arrange their generated colors into existing color buckets.

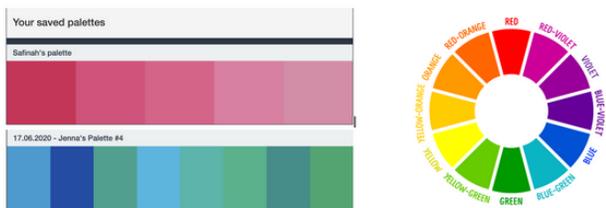


Figure 1: TryColor activity: students generate color palettes using a few starting colors (left), then they classify their generated colors into color wheel buckets (right).

Through this activity, students grasp the difference between algorithms that classify and those that generate. Students then learn about examples of AI systems that perform classification and generation and practice distinguishing between the two. Next, we introduce Generative Adversarial Networks (or GANS) as a type of AI that can generate media after being trained on datasets of media samples. Awareness of GANS in daily life is raised as students learn to distinguish GAN-produced media from human produced media and practice searching for tell-tale signs of algorithmically generated media in an activity called "GANs or not". We conclude the unit by connecting the core concepts to ethical considerations and career futures. After examining different works of art generated by GANS, we discuss the ethics of GANS (with prompts like "Who should get credit as the artist?"). Then we watch a video about an artist who uses AI as a tool for creative expression, and engaged students in ideating about how AI is utilized in artistic creation and what they would like to generate with GANS.

3.4.1 Conversion of Activities for Online Settings. We adapted the DAILY curriculum for remote synchronous online instruction in preparation for the online summer workshop. This entailed the conversion of several in-person classroom lessons to online activities. For instance, Unit 1 consisted of an *AI or Not* activity, in which students sorted physical cards with pictures of common technology (such as Alexa) as *AI* or *Not AI*. For the online adaptation, we made use of Google Drawing, and digital cards that students could drag into an *AI* side and a *Not AI* side. Each student was provided with their own copy of the Google Drawing with all the cards initially in the middle. They were tasked with sorting the cards individually into the two categories.

3.5 The AI Workshop

The AI workshop was offered in Summer 2020, in partnership with a youth serving community organization. Due to the COVID-19 pandemic, the course was taught virtually. Students met online, on Zoom, for three hours per day each weekday over two weeks. The course was taught by a team of researchers and educators. Each session typically started with the instructor introducing the unit's topic, followed by a whole-class activity, a small group or individual activity, a discussion relating to ethical implications, and connections to AI careers. Participants were grouped by age into three groups of 10 or 11 individuals for small group discussions and hands-on activities.

4 IMPLEMENTATION EXPERIENCES AND CHALLENGES

4.1 Student Participation and Engagement

Classroom observations revealed that the enactment of the DAILY curriculum was successful overall. Students completed all curricular activities outlined in Table 1 with high attendance rates (85 to 95% of all participants attended each day). They actively participated in almost all activities, including the technology explorations, interactive lecturing, embedded assessments, class discussions, daily reflections, and final projects. However, the percentages of students who submitted the completed work for each individual activity ranged from 80% to 30%. One reason for the varied submission

Table 1: The DAILY curriculum: interweaving of AI concepts, ethics, awareness, and careers

Unit	AI Concepts	Ethics	Awareness	Careers
Unit 1	What is AI?	Best PB&J and Ethical Matrix	AI or not AI?	
Unit 2	What are Decision Trees? Pasta-Land	Investigating Bias in AI	Examples of Bias in AI	What will your future career be?
Unit 3	What is Supervised Learning? Explore Teachable Machine	Algorithmic bias. Examples of algorithmic bias.	Supervised Learning in your life - AI Bingo activity	Inventory of me
Unit 4	What is a Neural Network? NN game	Where bias might be embedded in NNs.	Examples of where NN are used	Work personality and matched occupations
Unit 5	How to train a Supervised Learning model in Teachable Machine.	How accurate is your model? What is the potential for bias?	Image classification in your life	Planting seeds of STEM/AI careers
Unit 6	Classifier vs Generator. What are GANS?	Ethics of GANS	GANS or not	People working in AI fields. Creativity in STEM/AI careers
Unit 7	How do GANS work? Generator vs Discriminator game	Art & Ownership discussion	GANS in your life. Engage with various applications of GANS	AI's impact on my future field of work? People working in AI
Unit 8	What are Deepfakes?	Deepfakes and Misinformation	Deepfakes in your life. How misinformation spreads	People working in AI related fields continued.
Unit 9	What is Text Generation? Make a text generator	Best Ethics of text generation	Text generation in your life. Examples of text generation	People working in AI related fields continued.
Unit 10	Ethical Design of AI. Redesign YouTube	Ethical issues in AI (recap)	Stakeholders have different goals for AI	My Career Roadmap

rates may be related to the degree of interactivity with peers in the activities. For example, in the "Investigating bias" activity, students participated in impassioned discussions on examples of bias they noticed in everyday life, commented on each other's examples, and searched online to uncover new examples of biases (such as results of Google image searches). In contrast, submission rate for the corresponding assignment, that of documenting a search for bias, was low despite the high engagement of students.

Based on these findings, we posit that students preferred live exchange of their thoughts and ideas with peers to writing them down, when exploring bias and other ethics-related issues. In the "TryColor" activity, the low submission rate may be related to the gamified approach employed in the activity. Over half of the students spent almost all the allocated time on the game of generating their own colors and replicating others' colors. Many of them ended up with insufficient time to complete the assignment. In addition, we postulate that the low submission rate of assignments may be related to the Google Classroom we used to manage the enactment of the workshop. While this tool helped deliver the workshop activities in a well-structured way, it made the experience resemble school-based ones, which did not resonate well among youth in the out-of-school setting. It is understandable that students in this by-choice summer program felt reluctant to submit multiple assignments every day as they would in school.

4.2 Student Experiences with Activities

We examined students' daily reflections and interviews to reveal more detailed information about their experiences during the DAILY workshop.

4.2.1 Ethics in AI activities. The findings show that students mostly enjoyed ethics-related activities of Unanticipated Consequences

(Unit 8) and Investigating bias in AI (Unit 2). 10 out of the 11 students who volunteered to be interviewed after the workshop thought these activities were their favorites. The "Unanticipated Consequences of Technology" activity engaged students in working in small groups to first brainstorm one consequence of a technology (e.g., Lady GAN is a new pop star whose music is created completely by generated music and lyrics to match the style of Lady Gaga). The group then sent their consequence to other groups for the attendant consequences. Finally, the first group reviewed and reflected on the resulting consequences. Students expressed their excitement when working on this activity in the interview, e.g., *"That [activity] was hilarious because a lot of the consequences my group came up with totally turned the tables. Some of them made it much worse and some of them just made it much better, but it was total opposite from what the first consequence with Amelia [pseudonym of a student in the group] was."* During the post-interview, students were also able to internalize what they had experienced and to connect to ethical implications of technology design, *"because if people who make the AIs don't think enough through it, they can have hard consequences and bad consequences. People can get sued if they don't check over or people can get in like a lot of trouble and damage if they don't check over. But it [this activity] was funny, but it's realistic."*

4.2.2 Investigating bias. Similarly, students actively participated in the "Investigating bias in AI" activity where they learned about case studies of bias in AI systems, conducted investigations (searching images of physicists online, etc.) to explore bias present in common technologies, and discussed the implications. In the post-activity discussion, students were surprised to recognize that bias has been existent in many daily technologies, e.g., *"when we typically think of Google, we think it's objective, it's always right, but it's not always represented in the right way"*. We noticed that female students from

underrepresented groups were particularly active in the investigation and ensuing discussion. They talked about fairness, how to provide everyone with equitable instead of equal opportunities, and what sexism and stereotyping mean. In a classroom discussion, one female African-American student concluded that *“My takeaway is that AIs like Google and facial recognition, haven’t been updated in a long time because you can clearly see the bias in that. And even though we think that these things are very... very smart, but I never understood or considered that they might have been really bias until going over these things.”* These findings are consistent with the literature that women and underrepresented minorities are often interested in issues that address societal and ethical concerns [37].

4.2.3 Design and build experiences. The technology exploration and creation activities supported students in making sense of the underlying AI concepts. Students in total spent approximately five hours on a series of activities using Google’s Teachable Machine (in Units 3 and 5). They first learned and explored how to use the tool to train a supervised learning model to recognize faces, gestures, and voices. Then they experimented with training a model on a biased dataset (one with more pictures of dogs than cats), observed the results, and brainstormed and attempted to fix the bias. As a final activity they explored integrating the supervised learning model they trained using Teachable Machine into Scratch animations (e.g., when the model recognized certain voices, the sprite would wave hands). Our observations showed that although students had challenges in the last activity due to limited Scratch skills or computing power of their computers, students in general developed a solid understanding of supervised learning and how to mitigate potential biases in supervised learning systems. Further this experience of training an AI model enabled some students to develop a sense of control over technologies as one student reflected in their post-interview, *“I feel like we could do it,...so teachable machine is, we’re in control, but we’re teaching the teachable machine how to do something.”* This sense of agency or control over technology has been long recognized as a key factor in impacting peoples’ interactions with and perceptions of technology [27]. It may have helped the participating youth become more interested in and less anxious about AI.

4.2.4 Participatory simulations. Leveraging a participatory simulation approach [24], the Neural Network game engaged students in playing the role of nodes in training of a network. Students captioned a mystery image by passing on information through different “layers” of the network and experienced how nodes, links, and layers worked to train the network. Students found the analogy between feed-forward mechanism and the “telephone game” helpful in helping make sense of the progression of information.

4.2.5 Generative AI activities. Activities around generative AI were engaging yet difficult for students. For instance, the “colors” activity (see an earlier description of the detailed activity) was frustrating as one student mentioned in her interview *“because there was a little too much of one and too much of the other. Even if you got like, at least like a 96% close to it, for me, it still didn’t allow me to pass... Sometimes I did give up.”* Despite the frustration, the student still found it interesting because she could *“see how people would mix up the colors and all.”* Another example is the “Deepfakes” activity in

which students reviewed videos, images, and texts and determined whether they were generated by AI. Students were astonished to find out that all the files were generated by AI. Many students admitted that even after receiving instructions on how to detect Deepfakes, they still had difficulties.

4.2.6 Career development activities. The career training sessions immersed students in exploring their future jobs. In the “Inventory of Me” activity students discovered their work personalities and found jobs that match them. Students were surprised to find their matched job options as sometimes they had never heard of them. The exploration of how AI has been and may be adopted in different jobs sparked their interest in AI. In the post-interview, one girl thought that AI will impact her future job *“because when we were doing the job career search, a lot of the jobs that popped up were jobs that I was interested in and it was connected to AIs. So I feel like that definitely shows that a job that I want will be connected to AI in some way.”*

4.3 Implementation Challenges

Remote synchronous online instruction due to COVID-19 presented several implementation challenges. In online versions of unplugged activities, the benefits of working with physical materials and embedding activities in a physical space were lost. For example, in the “PastaLand” activity, students’ ability to identify key features of the pasta was limited by the graphical representation we chose. The relative sizes of the different pasta as well as the textures could not be accurately portrayed. Similarly, the online environment limited the use of a project-based learning approach. The difficulties encountered with supporting students’ collaboration online limited opportunities to help students apply theoretical knowledge to constructing tangible artifacts.

The online environment also shifted interactions with peers and instructors from semi-private to public. For instance, in an online environment, we found that collaboration while debugging moved from an intimate (semi-private) interaction to an interaction shared with the entire classroom. Discussions were central to several of our activities since the curriculum focused on not just teaching AI concepts but also discussing ethical issues associated with AI. Lack of access to a microphone or video camera, or background noise hindered effortless conversation in these discussions, and at times we had to rely on chat. Student engagement was challenging to maintain and track, since unlike in-classroom learning, we had limited access to visual cues and emotional responses from students.

Another implementation challenge pertained to discussing algorithmic bias, as well as its implications in AI, and its societal implications such as unanticipated discrimination against underrepresented minority groups with students from these groups. While central to the curriculum, we aimed to present these concepts with sensitivity to the audience. With guidance from the community partner, we discussed how the AI industry was seeking to address these biases and how they, as future AI engineers, can help to identify and mitigate some of these biases. While it is important to discuss the existence and implications of algorithmic bias, we refrained from framing the discussions of bias with students from under-represented groups such that we put the onus on them for fixing a problem in AI that they did not create. Hence, instead of

framing the discussion as how they can solve the problem of bias in AI, we framed it as how the AI community can solve it. Lastly, while we discuss several negative implications of AI technologies such as bias, generation of fake media, and misinformation, we notice a need to include more positive examples of AI technologies, such as applications in healthcare, education, and art.

5 LEARNING OUTCOMES

The learning outcomes were mainly evaluated using two instruments administered before and after the workshop: (1) the AI Concept Inventory that assesses youths' knowledge and skills in AI through three validated sub-scales: AI general concepts, logic systems, and machine learning, and (2) the Attitudes toward AI and AI careers survey that explores youths' interest in AI and related careers, AI career awareness, and career adaptability. All questions were drawn and modified from validated instruments, e.g., the career adaptability includes 10 items based on the revised Career Futures Inventory [35, 41]. Detailed information of the two instruments can be found in Lee et al. 2021 (in press) [?].

We found a statistically significant increase in youths' AI concept inventory scores from pre- to post-test ($t(17)=-2.09$, $p<.05$, $d=.54$). The subscale with the largest increase in scores was machine learning (including both unsupervised and supervised learning). Students on average had a score increase of 1.58—meaning that over half of the students correctly answered at least two more questions about machine learning concepts and processes on the post-test. We also found significant increase in AI career awareness ($t(18)=3.58$, $p<.01$, $d=1.00$) and career adaptability skills ($t(18)=1.013$, $p<.001$, $d=1.40$) after the workshop. Students also increased their interest in AI after the intervention, yet the increase was not statistically significant. One reason may be the ceiling effect. Participating youth were highly interested in AI before the workshop (Mean=3.59 out of 5) and maintained their interest level on the post-test (Mean=3.71).

6 CONCLUSIONS

Preparing youth to become knowledgeable citizens and workers in the age of AI is essential for training the next generation of the workforce. This paper reports the design and implementation of an AI workshop among middle school aged youth because middle school years are an important time to develop their interest in and shape their views of AI and related jobs. This workshop aims to develop AI literacy through an integration of AI concepts, ethical and societal implications of AI, and the adoption of AI in future jobs. The workshop also aims to provide a low-barrier to entry for learners with no pre-requisite computing resources or knowledge of computing systems, thereby enabling learners with no CS background to gain a foundational knowledge of AI concepts. We pilot tested this workshop among 31 middle school youths from groups underrepresented in STEM and Computing in a self-selected summer program. We found significant increases in students' conceptual understanding of AI, knowledge of potential bias of AI, and ability to adapt to future AI-empowered jobs. This indicates that the approach of combining AI, ethics, and careers to develop AI literacy is feasible and suitable for middle school youth.

Consistent with previous AI ethics curricula, we have found middle school students were highly engaged in discussions of ethics,

in the context of fairness and bias in AI. Students were capable of reflecting on their own behaviors and its impacts on others, and can extend this reflection to encompass how various stakeholders in AI systems can design systems based on their values and needs. Students could identify how and why AI applications can be unfair to themselves and others. For instance, in the post-test interview, one student said, *"I think that bias [in AI] could be harmful for us because a lot of people aren't fairly represented, so they wouldn't be able to use it as much."*

Further, we found that the youths talked with their family members about AI. A student in his post-test exit interview told us that *"I asked my mom that day when we were talking about [...] if AI could probably take our jobs. I asked her, 'Are you scared if AI takes your job?' She's like, 'Yes, but I know I'll have another job.' So, it's really cool what AI could do and make new things for other people."* Through the conversation with his mom, the student was active in leveraging existing familial and aspirational capital [17] to make sense of the potential impact of AI and to inform and refine his existing understanding and perceptions of AI. Meanwhile, this quote provides further evidence that our approach may have sparked high interest among students such that AI arises during family conversations. Research on science learning has suggested that conversations about science topics are not typical occurrences for all youth, especially for those from underrepresented groups in STEM [18, 34].

7 FUTURE WORK

Based on our pilots and findings, we aim to further revise our curriculum and assessment methods. First, we plan to add more discussions around positive application areas of AI such as healthcare and education. We will frame the mitigation of bias in AI conversation such that the onus of fixing the problem does not lie on students from under represented minority groups. Second, we aim to include more project-based activities where students can apply their own interests, and ethical considerations to an AI-enabled project. For instance, we plan to include a generative modeling activity, where students pick their own training datasets, and then generate and share synthetic data such as text, visual art or music, while reflecting on the ethical implications of the generative media they choose. This may necessitate being co-located in classrooms and having computing capability to train large models. Another final project could entail redesigning recommender systems while keeping in mind stakeholders and their values, and checking for inherent biases in the system. Lastly, we aim to scale this curriculum to teachers and students across the country for wider access. We aim to work with educators to explore the efficacy of the DAILY curriculum both in formal middle school curricula and informal settings, such as after-school or summer programs. Once we collect and analyze data from the additional deployments of the DAILY curriculum, we plan to develop an AI literacy competency model to further contribute to the field of AI Education.

8 ACKNOWLEDGMENTS

This work was funded by the National Science Foundation under grant #2022502. We are grateful to the teachers, facilitators and

students participating in middle school STEM programs in Massachusetts. We thank Grace Kim, Carolyn Song, Jenna Hong and Yihong Cheng for assisting us with data collection and analysis.

REFERENCES

- [1] 2018. AI4k12.org. <https://github.com/touretzkyds/ai4k12/wiki>
- [2] 2020. The DAiLy Curriculum. <https://aieducation.mit.edu/daily/>
- [3] 2020. Machine Learning for Kids. <https://machinelearningforkids.co.uk>
- [4] 2020. TryColors. <https://trycolors.com/>
- [5] 2020. Young Teens (12-14 years of age). *Center for Disease Control and Prevention* (2020). <https://www.cdc.gov/ncbddd/childdevelopment/positiveparenting/adolescence.html>
- [6] Safinah Ali, Hae Won Park, and Cynthia Breazeal. 2020. Can Children Emulate a Robotic Non-Player Character's Figural Creativity?. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. 499–509.
- [7] Safinah Ali, Blakeley H. Payne, Randi Williams, Hae Won Park, and Cynthia Breazeal. 2019. Constructionism, Ethics, and Creativity: Developing Primary and Middle School Artificial Intelligence Education. In *International Workshop on Education in Artificial Intelligence K-12 (EDUAI'19)*.
- [8] Louise Archer, Jonathan Osborne, Jennifer DeWitt, Justin Dillon, Billy Wong, and Beatrice Willis. 2013. ASPIRES: Young People's Science and Career Aspirations. *Age 10* (2013), 14.
- [9] David L Blustein. 2019. *The importance of work in an age of uncertainty: The eroding work experience in America*. Oxford University Press.
- [10] Andrew Perrin Brooke Auxier, Monica Anderson and Erica Turner. 2020. Parenting Children in the Age of Screens. *Pew Research Center* (2020). <https://www.pewresearch.org/internet/2020/07/28/childrens-engagement-with-digital-devices-screen-time/>
- [11] Katherine P Dabney, Robert H Tai, John T Almarode, Jaimie L Miller-Friedmann, Gerhard Sonnert, Philip M Sadler, and Zahra Hazari. 2012. Out-of-school time science activities and their association with career interest in STEM. *International Journal of Science Education, Part B 2*, 1 (2012), 63–79.
- [12] Susanne A Denham, TM Wyatt, Hideko Hamada Bassett, D Echeverria, and SS Knox. 2009. Assessing social-emotional development in children from a longitudinal perspective. *Journal of Epidemiology & Community Health* 63, Suppl 1 (2009), i37–i52.
- [13] Matthew J Drake, Paul M Griffin, Robert Kirkman, and Julie L Swann. 2005. Engineering ethical curricula: Assessment and comparison of two approaches. *Journal of Engineering Education* 94, 2 (2005), 223–231.
- [14] Stefania Druga. 2018. *Growing up with AI: Cognimates: from coding to teaching machines*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [15] Stefania Druga, Sarah T Vu, Eesh Likhith, and Tammy Qiu. 2019. Inclusive AI literacy for kids around the world. In *Proceedings of FabLearn 2019*. 104–111.
- [16] J Paul Gibson. 2012. Teaching graph algorithms to children of all ages. In *Proceedings of the 17th ACM annual conference on Innovation and technology in computer science education*. 34–39.
- [17] Habig B Gupta P and Adams JD. 2020. Disrupting deficit narratives in informal science education: Applying community cultural wealth theory to youth learning and engagement.(in press). (2020).
- [18] Shirley Brice Heath. 1986. Beyond language: Social and cultural factors in schooling language minority students. *CA: California State Department of Education* (1986).
- [19] Tom Hitron, Yoav Orlev, Iddo Wald, Ariel Shamir, Hadas Erel, and Oren Zuckerman. 2019. Can Children Understand Machine Learning Concepts? The Effect of Uncovering Black Boxes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [20] Ken Kahn and Niall Winters. 2017. Child-friendly programming interfaces to AI cloud services. In *European Conference on Technology Enhanced Learning*. Springer, 566–570.
- [21] Minsuk Kahng, Nikhil Thorat, Duen Horng Polo Chau, Fernanda B Viégas, and Martin Wattenberg. 2018. Gan lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 1–11.
- [22] Neil K. Kaneshiro. 2018. Adolescent development. *MedlinePlus* (2018). <https://medlineplus.gov/ency/article/002003.htm>
- [23] Neil K. Kaneshiro. 2018. School-age children development. *MedlinePlus* (2018). <https://medlineplus.gov/ency/article/002017.htm>
- [24] Eric Klopfer, Susan Yoon, and Judy Perry. 2005. Using palm technology in participatory simulations of complex systems: A new take on ubiquitous and accessible mobile computing. *Journal of Science Education and Technology* 14, 3 (2005), 285–297.
- [25] Robert W Lent, Steven D Brown, and Gail Hackett. 2002. Social cognitive career theory. *Career choice and development* 4 (2002), 255–311.
- [26] Jing Li, Wen Wang, Jing Yu, and Liqi Zhu. 2016. Young children's development of fairness preference. *Frontiers in psychology* 7 (2016), 1274.
- [27] Hannah Limerick, David Coyle, and James W Moore. 2014. The experience of agency in human-computer interactions: a review. *Frontiers in human neuroscience* 8 (2014), 643.
- [28] 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.
- [29] Adam V Maltese and Robert H Tai. 2010. Eyeballs in the fridge: Sources of early interest in science. *International Journal of Science Education* 32, 5 (2010), 669–685.
- [30] Jack Nicas. 2018. How YouTube Drives People to the Internet's Darkest Corners. *The Wall Street Journal* (February 2018).
- [31] Barbara M Olds and Ronald L Miller. 2004. The effect of a first-year integrated engineering curriculum on graduation rates and student satisfaction: A longitudinal study. *Journal of Engineering Education* 93, 1 (2004), 23–35.
- [32] B.H. Payne. 2020. *Can my algorithm be my opinion?: An AI + Ethics Curriculum for Middle School Students*. Master's thesis. Massachusetts Institute of Technology, Media Lab, Cambridge, MA, USA.
- [33] Jean Piaget. 2014. *Studies in reflecting abstraction*. Psychology Press.
- [34] Barbara Rogoff et al. 2003. *The cultural nature of human development*. Oxford university press.
- [35] Patrick J Rottinghaus, Kristine L Buelow, Anna Matyja, and Madalyn R Schneider. 2012. The career futures inventory—revised: Measuring dimensions of career adaptability. *Journal of Career Assessment* 20, 2 (2012), 123–139.
- [36] Mark L Savickas. 1997. Career adaptability: An integrative construct for life-span, life-space theory. *The career development quarterly* 45, 3 (1997), 247–259.
- [37] Susan C Seymour and Susan Christine Seymour. 1999. *Women, family, and child care in India: A world in transition*. Cambridge University Press.
- [38] Michael Skirpan, Nathan Beard, Srinjita Bhaduri, Casey Fiesler, and Tom Yeh. 2018. Ethics education in context: A case study of novel ethics activities for the CS classroom. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education*. 940–945.
- [39] Elisabeth Sulmont, Elizabeth Patitsas, and Jeremy R Cooperstock. 2019. Can You Teach Me To Machine Learn?. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*. 948–954.
- [40] Maciej M Syslo and Anna Beata Kwiatkowska. 2014. Playing with computing at a children's university. In *Proceedings of the 9th Workshop in Primary and Secondary Computing Education*. 104–107.
- [41] CFIR Research Team et al. 2017. Consolidated Framework for Implementation Research (CFIR).
- [42] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What should every child know about AI?. In *Proceedings of the AAI Conference on Artificial Intelligence*, Vol. 33. 9795–9799.
- [43] Randi Williams, Hae Won Park, and Cynthia Breazeal. 2019. A is for artificial intelligence: the impact of artificial intelligence activities on young children's perceptions of robots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [44] Randi Williams, Hae Won Park, Lauren Oh, and Cynthia Breazeal. 2019. Popbots: Designing an artificial intelligence curriculum for early childhood education. In *Proceedings of the AAI Conference on Artificial Intelligence*, Vol. 33. 9729–9736.
- [45] Chip Wood. 1997. *Yardsticks: Children in the Classroom Ages 4-14. A Resource for Parents and Teachers*. ERIC.
- [46] Abigail Zimmermann-Niefield, Makenna Turner, Bridget Murphy, Shaun K Kane, and R Benjamin Shapiro. 2019. Youth learning machine learning through building models of athletic moves. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*. 121–132.