

Teacher Perspectives on How To Train Your Robot A Middle School AI and Ethics Curriculum

Randi Williams, Stephen P. Kaputsos, Cynthia Breazeal

MIT Media Lab
Cambridge, MA 02142
{randiw12, kaputsos, cynthiab}@media.mit.edu

Abstract

To enable a diverse citizenry to fully participate in future society, we must prepare all students to construct and critique emerging technologies like Artificial Intelligence (AI). Classrooms are important spaces to teach students these skills, however there are few AI curricula that have been developed for and used by K-12 teachers. We developed the *How to Train Your Robot: AI and Ethics Curriculum* for middle school teachers who want to introduce AI to their students. This paper describes the curriculum and professional development we used to prepare teachers to run a five-day AI course. Before and after they ran the curriculum, we interviewed teachers to understand their opinions on pedagogical approaches to teaching AI, meeting students' needs, and the feasibility of doing the activities in the classroom. Our results indicate that, with appropriate training, even teachers who were new to computer science felt prepared and successfully engaged their students in the topic. We hope our insights will inform future efforts to realize AI education in primary and secondary classrooms.

Introduction

Artificial Intelligence (AI) impacts multiple areas of today's society, including how people work, learn, govern, and entertain themselves. A lack of skills in emerging technologies could limit a person's ability to actively participate as a citizen and compete for the highest paying jobs (Tuomi 2018). Thus, many arguments in support of AI education in primary and secondary school have come forth in recent years (Zimmerman 2018; Touretzky et al. 2019). Going one step further, we believe that AI education efforts must target K-12 classrooms at public schools to make knowledge AI accessible to the most diverse group of students.

We developed the *How to Train Your Robot* curriculum for teachers who want to explore AI and ethics with their middle school students (ages 11-14). We conducted two rounds of classroom studies with in-service middle school teachers to get teachers' feedback on the activities. The first round of studies occurred in classrooms over five full school days. Then, due to the health emergency caused by an outbreak of SARS-Cov-2, we converted the curriculum and ran it as a synchronous, online course for five, 2.5-hour sessions.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Background

AI Education

Published studies on AI curricula in K-12 classrooms date back to at least 2010, with the majority of the work being published in the past 3 years (Heinze, Haase, and Higgins 2010; Long and Magerko 2020). In the United States of America, the AI4K12 initiative outlined national guidelines on what every K-12 student should understand about AI. The guidelines describe the Five Big Ideas of AI: how AI enables computers to perceive their environment, represent and reason about knowledge, learn from data, interact naturally, and impact society (Touretzky et al. 2019). Our curriculum focuses on Big Ideas #3: that computers can learn from data and #5 that AI can impact society in both positive and negative ways (Touretzky et al. 2019). Lao and Long et al. expand on this work by providing AI competency frameworks (Lao 2020; Long and Magerko 2020). With regard to these frameworks, we also teach competencies 1 and 2, how to define intelligence and recognize AI (Long and Magerko 2020), and skills 2 and 3, how to plan a project and create machine learning artifacts (Lao 2020).

Machine Learning and Robotics in K-12 AI Education

Many AI interventions use programming activities enable students to learn about AI by tinkering (Lane 2021; Druga 2018; Hitron et al. 2018; Kahn et al. 2018; Hitron et al. 2019; Tang et al. 2019; Zimmermann-Niefield et al. 2019; Bhatia and Lao 2020; Carney et al. 2020). Our curriculum leverages tools from the Machine Learning for Kids and Teachable Machine projects so that students can learn through making (Long and Magerko 2020). Interventions also use non-programming activities to help students engage with AI in a variety of contexts (TechGirlz 2018; Vahrenhold, Cutts, and Falkner 2019; DiPaola, Payne, and Breazeal 2020; Payne 2020). We used Payne's metaphors, such as comparing algorithms to recipes, and non-programming activities, such as building a biased cat-dog classifier, to help students gain new perspectives on AI before they programmed it (Payne 2019).

Although robotics education is not our focus, we integrated a physical robot into the course to promote student engagement through embodied interaction (McConnell

1996; Kumar 2004; McNally and Klassner 2007; Koski, Kurhila, and Pasanen 2008; Talaga and Oh 2009; Touretzky 2014; Long and Magerko 2020). Related fields such as mathematics and computer science education attest that physical manipulatives make ideas more concrete and, thus, easier to grasp (Papert 1980; National Council of Supervisors of Mathematics 2013). Researchers saw these same benefits when they included robotics in their AI classrooms. Kumar et al. used LEGO robots in an undergraduate AI course to physically situate the AI algorithms, offering another medium to understand them beyond code and equations. The drawbacks of using physical robots to teach AI are that robots are time-consuming to physically maintain and do not always behave as expected (Kumar 2004; Koski, Kurhila, and Pasanen 2008; Talaga and Oh 2009). Therefore, educators must find a way to balance the amount of time students spend working with the robots and actually learning about AI (Kumar 2004; Koski, Kurhila, and Pasanen 2008; Talaga and Oh 2009).

AI + Ethics Education

Ethics is an extremely important, yet under-taught topic in many Computer Science and AI classes (Payne 2020; Fiesler, Garrett, and Beard 2020). In the MIT AI and Ethics Middle School Curriculum developed by Payne, technical concepts are taught in situ with ethical ones. This promotes student awareness of both the ethical implications of existing AI technologies and the importance of designing technology from an ethical foundation (DiPaola, Payne, and Breazeal 2020; Payne 2020). Our curriculum leverages many of the tools and activities developed by Payne to help students develop this foundation.

Teacher Training

The primary goal of this work was to enable practicing teachers to implement a week-long AI course in their classroom. Although many have recognized the importance of including AI education in schools, there has been little work evaluating interventions in classrooms and engaging with K-12 teachers (De La Higuera 2019; Vazhayil et al. 2019; Marques, Gresse von Wangenheim, and Hauck 2020). Vazhayil et al. ran a study where they trained 34 secondary school teachers in India to teach a project-based AI curriculum based on Machine Learning for Kids (Vazhayil et al. 2019). They interviewed teachers after training them and found that, although teachers were eager to begin using the tools, there were concerns about how to efficiently run an exploratory, project-based course with students. Our work builds on that paper by looking at teacher perspectives after they were trained and ran the curriculum themselves.

Curriculum Design

The *How to Train Your Robot* curriculum includes activities, hardware and software tools, and teacher training materials. In designing the curriculum, we prioritized hands-on activities, accommodating novice teachers and students, real-world relevance, and cost-effectiveness.

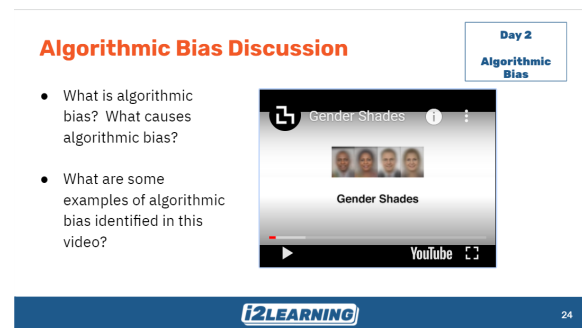


Figure 1: This is a slide we used for the Algorithmic Bias Discussion activity. Teachers play the video then lead the class in answering discussion questions.

Activities

Every day, students explored different ethical and technical concepts in AI and machine learning. In the ethics modules, they learned about the positive and negative impacts of AI, product design trade-offs, and how to conduct stakeholder analysis using ethical matrices. Many of these activities are adaptations of AI and ethics activities designed for middle school students by Payne (Payne 2019). In the technical modules, students learned about text and image classification, algorithmic bias, and how to build machine learning models. At the end of the week, students completed a final project and presented them at a showcase for their family and friends. In their projects, students applied their technical and ethical skills as they prototyped and critiqued their AI inventions. Table 1 outlines the major activities for each day. Full descriptions of the activities can be found in the Appendix¹.

We made the curriculum accessible to teachers and students who were new to computer science by requiring minimal prerequisites and including some non-programming activities. We did not require or expect students or teachers to have any prior programming or computer science experience. We taught teachers all content knowledge they needed to know and built in tutorials to support students. Also, half of the activities in the original in-person version of the curriculum were unplugged; in the online version of the class we describe them as “non-programming.” Balancing unplugged and plugged activities allowed students to exercise and develop their critical thinking as well as computational thinking skills.

Hardware and Software

How to Train Your Robot uses a custom build of the Scratch programming blocks². Scratch is designed to be welcoming to students who are new to programming, yet powerful enough to interest students with prior experience (Resnick et al. 2009). We built extensions that integrate custom ma-

¹The appendix for this paper, with descriptions and timing of activities, can be found at <https://github.com/mitmedialab/prg-extension-boilerplate/blob/robotafe/HTTYRCourseOverview.pdf>

²LLK Scratch Blocks on Github, <https://github.com/LLK/>

Session	Activity	Learning Goals (Students will be able to...)
Session 1	AI or Not Ethical Dilemmas Video Intro to Scratch & Robots	Define AI, reason about what makes something AI or Not Employ strategies to reason through ethical dilemmas Use block-based programming to complete mini-projects with robots
Session 2	Alexa Pizza Delivery App Image Classification Algorithmic Bias Videos & Articles Teachable Machine + Scratch	Define algorithms and design an algorithm for real world use Curate datasets and use them to train image classifiers, understand neural networks Discuss the implications of algorithmic bias and what can be done to mitigate it Build custom image classification models to program robots in Scratch
Session 3	Ethical Matrix: Redesign Alexa Exploring Word Analogies Text Classification + Scratch	Use stakeholder analysis to unpack the implications of technology design Understand word vectors and how they encode language Discuss bias in large datasets and how to mitigate it Explain the k Nearest Neighbor algorithm and how it is used in machine learning Build custom text classification models to program robots in Scratch
Session 4	Final Project Brainstorming Final Project Planning Final Project Work Time	Do research to generate ideas for final projects Use project planners and ethical matrices to develop final project ideas Employ time management to work on an open-ended project Give and receive peer feedback
Session 5	Showcase	Complete a final project that applies a lesson from the week Create a presentation to explain final project to a general audience

Table 1: *How to Train Your Robot* activities and learning goals by day.

chine learning models, speech-to-text, text-to-speech, and micro-controller programming blocks. In the first set of studies, we used ScratchX and in the second set we upgraded to Scratch 3.0, the most recent version of Scratch (see Figure 2).

To build and create machine learning models, we wanted to make sure students had access to free, browser-based tools that they could access even outside of our course. In the first set of studies we used Machine Learning for Kids³, which is a kid-friendly interface for making image, text, audio, and number classification models. To use this tool, teachers must register for a developer account through IBM Watson’s cloud API.

Due to limitations with using a cognitive service’s API in classrooms, which has also been noted in other AI education studies (Vazhayil et al. 2019), we moved to Google’s Teachable Machine and our own Scratch Text Classifier tool for the second study. With Teachable Machine, students can train and run inference on image recognition models in their browser - meaning that data does not leave students’ computers (Carney et al. 2020). We developed an extension for the Scratch interface that allows students to export their models and use them in Scratch programs (Jordan et al. 2021). For text classification, we built an interface that lets students train and run inference models directly in the Scratch interface (Reddy, Williams, and Breazeal 2021). More information about these tools can be found in our

GitHub repository⁴.

The curriculum includes a robot because we wanted to increase student engagement and encourage students to bridge their knowledge across physical and digital spaces (Papert 1980; National Council of Supervisors of Mathematics 2013; Kumar 2004; Koski, Kurhila, and Pasanen 2008; Talaga and Oh 2009). Each robot cost less than \$50 and connected to students’ Chromebooks and laptops using Bluetooth. We intentionally designed the technology to be low cost and easy to set up so that it would be feasible for teachers to run it themselves. In the first study, we hand built Arduino-based robots and provided them to schools. In the second round, improvements to the Chrome browser’s web Bluetooth functionality allowed us to upgrade to micro:bit robots which are sold by a company⁵. Both robots include line sensors, ultrasonic distance sensors, motors, and RGB LEDs. The micro:bit robots also include a piezo buzzer, two push buttons, and a 25-LED display. Both robots are shown in Figure 2.

Teacher Training Materials

We provided teachers with an educator guide, slide deck, student worksheet materials, programming guides, and synchronous training sessions. Written by AI experts and a middle school teacher, the educator guide provided teachers with comprehensive instructions on how to run activities includ-

³Machine Learning for Kids by Dale Lane, <https://machinelearningforkids.co.uk/>

⁴PRG AI Blocks on Github, <https://github.com/mitmedialab/prg-extension-boilerplate/tree/robotafe>

⁵Yahboom Tinybit, <https://category.yahboom.net/products/tinybit>




Figure 2: Scratch interfaces and programmable robots from both studies. The top shows the technology from our initial study, Arduino robots and the ScratchX interface. The bottom shows the tools from our second study, micro:bit robot and the Scratch 3.0 interface.

ing discussions and programming. Lesson plans employ the universal design for learning framework and provide variations of activities and links to the real world to help teachers connect the subject matter to their students' strengths and interests (Rose and Meyer 2006; Heinze, Haase, and Higgins 2010).

AI researchers led the teacher training, conducting activities as though teachers were students to help them learn the pedagogical tools they would use in class. During training, researchers focused on helping teachers navigate the hardware and trickier material in the curriculum: articulating a clear definition of AI and developing a balanced view of AI as potentially helpful and harmful.

Initial Study

Participants

In the first iteration, we recruited seven teachers to run the curriculum as a part of a statewide STEM initiative. Of those seven, three participated in the research portion. One teacher came from a Title 1 school in an urban area (P1). The teacher taught Math and was comfortable with technology but had never taught computer science before. The other two teachers came from a small school (<100 students) in a rural district. One teacher was trained at a technical school and was comfortable with programming (P2). The other teacher

was beginning their first year teaching science and math and felt less comfortable with technology (P3). Students in both schools had previous experience with Scratch.

Procedures

Teachers underwent a two-day training program with researchers, a condensed version of the actual class. One problem we had in the first training session was that the robots were not ready beforehand so teachers used a similar alternative, a larger Arduino robot that connected to computers via USB.

When teachers ran the curriculum in their classrooms, they had complete freedom to adjust activities as they saw fit. They sometimes got in touch with the curriculum designers for help with technical problems. Their requests were centered around the first day, when first connecting robots to computers, and the fourth day, after a server that supported one of the programming activities was temporarily disabled.

At the end of the week, we asked teachers the following questions:

1. Did doing the AI curriculum change how your students think about AI?
2. How do you feel the week went for you as the teacher? What are some low-lights and highlights?
3. How effective was the training workshop? Was there any further training that would have been useful?
4. What are the most important lessons about this topic that students should walk away with?

To gather insights from teachers' answers, two curriculum designers (one authored this paper) reviewed them to identify common responses. A limitation of this procedure is that it was done in the spirit of designing the next iteration of the curriculum, not to systematically analyze teachers' responses for research purposes.

Results

In teachers' interviews, we observed prominent themes around student engagement, teacher preparation, and the use of robots in the classroom. Overall, students were clearly engaged in learning about AI, however there was some nervousness from those who were newer to the topic.

"They all have background knowledge of AI, so they are engaged." "Even kids that aren't engaged in Math, they were very excited." (P1)

"Everybody was more engaged than usual." "Some students sat back and let others do most of the work." (P2)

The unplugged activities and discussions received mixed responses. While some teachers embraced the nuances of the "AI or Not" and ethics discussions, others found them difficult to lead.

"The [AI or Not activity], the kids got stuff out of." "[Ethical dilemmas was] more confusing for students." "[Algorithmic bias was an] uncomfortable conversation, but kids understood it." (P1)

"Doing the [AI or Not] activity made them realize how much AI is involved in their life." (P2)

Finally, we saw that the hardware caused extra difficulty for teachers. But teachers felt it might be worth it because students found it exciting and built important skills like resilience.

“The robot was the biggest issue.” (P1)

“They used the same skills they use every day but they were into it because there were robots.” “The kids got more into the experience of robot and programming than diving deep with the AI.” (P2)

“[I saw them get] better at helping each other. Teamwork, perseverance, growth mindset.” “Teachers and kids had to do a lot of problem-solving. Some of the students could handle it, but others had a harder time.” (P3)

Using teachers’ feedback, we revised the activities and the tools. We made the curriculum more clearly align with middle school reading, writing, and CS standards and gave teachers more options to approach activities for students with different interests and strengths. We also completely redesigned the technology for easier setup. Machine Learning for Kids limits the number of models students can make and ScratchX required students to manually add an extension to control the robot, which was clunky and difficult. The Teachable Machine and Scratch Text Classifier extensions allowed students to make unlimited models. Plus, moving to a Scratch 3.0 repository form allowed us to streamline the experience of loading different extensions. Finally, we replaced the hand-built Arduino robots with more reliable, commercial robots and further simplified the setup procedure.

Second Study

Participants

In the second iteration of teacher training, participants included 7 teachers (P01-P07) who all taught Computer Science or STEM classes. Teachers had very diverse backgrounds in terms of previous teaching experiences, with a range of teaching experience from 2-24 years and an average of 12 years. All teachers work in Title 1 public schools, except for one teacher who directs a home school network that includes students with special needs (P06). Most teachers and students were new to the subject of AI except for two: one who had previously explored tools like Teachable Machine in their middle school classroom (P03) and another who had discussed AI topics with high school students before (P02).

Procedures

Due to the global health crisis caused by an outbreak of SARS-Cov-2, we moved the second set of training and workshops online. Each workshop consisted of five, 2.5-hour sessions; we shortened the time from full-day lessons so that students did not have to endure long video calls. More information about the timing and content of activities can be found in the Appendix⁶.

⁶Find the appendix at <https://github.com/mitmedialab/prg-extension-boilerplate/blob/robotafe/HTTYRCourseOverview.pdf>

Rather than working within their classrooms, teachers recruited six students to participate in the online course co-taught the class with the other teachers and two researchers. Researchers and teachers were equally involved in delivering online instruction. Researchers delivered 10-minute introductions to concepts each day, organized transitions between activities, and provided support during activities. Pairs of teachers worked together to lead the activities with small groups of students. Every day after class, researchers spent one hour with teachers preparing them for the next day of class.

At the end of every day we asked teachers:

1. Which activities were most engaging for your students?
2. What were some things that your students struggled with?

At the end of the week, we asked teachers:

1. Did the material in this course change your opinion about AI or teaching AI to students in any way?
2. How engaged were students in the course material?
3. What were the most important skills and ideas that your students learned in this course?
4. What is something new that you bring to this course or would like to add to it?
5. What parts of this course would you bring to your classroom and what would you leave behind?

To analyze teacher responses, two researchers (both authors of this paper), independently reviewed teachers’ answers and categorized quotes into one of five categories: technology, training, knowledge gain/insight, teaching strategy, and student engagement. Then, researchers went through the quotes to inductively identify outstanding themes.

Results

Student Learning and Engagement

Teachers expressed that the curriculum effectively engaged students throughout the five sessions. Some students even explored content beyond class hours.

“I was impressed with how they were collaborating...That even carried over into after hours. They asked questions on [Google] classroom and stuff.” (P02)

“They were asking ‘can I do more?’ It’s nice to see that they wanted to go beyond the scope.” (P03)

“They were all really into what we were doing.” (P05)

Student engagement expressed itself through students’ perseverance in moments of difficulty.

“They were like ‘I want this to work.’ and they tried really hard.” (P04)

“[They] asked me for help in a ‘phone a friend’ way looking for advice not all of the answers...By the time I figured it out they had already figured it out because they kept trying” (P05).

Teachers developed new perspectives regarding AI technologies.




Figure 3: Each day, we asked our teachers which activities were most engaging. The x-axis is a count of how many teachers mentioned the activity (they could mention more than one) and the y-axis refers to the session a specific activity occurred. Activities that teachers did not mention are in light gray.

“I’d never looked at the technology that way or thought about it that way” (P01)
 “It helped me understand it better so that I can break it down [for my students].” (P04)

Teachers constantly mentioned the programming activities as being the most engaging (see Figure 3). They attributed students’ understanding to their ability to participate in all parts of the AI development process - training and testing models, then utilizing them in programs.

“You can use your model combined with programming to make everything into a project. It puts all of the pieces into one.” (P03)

Students’ knowledge also manifested in their final projects. Students were encouraged to apply what they had learned about machine learning to an area or problem they cared about. Their work, detailed more thoroughly in our other paper (Williams 2021), ranged from educational technologies to scientific tools and mostly benefited the students themselves, their friends or families, or other kids.

“[The] projects at the end really gave an insight into their personal values and how they perceive the world.”
 “They were clearly applying technical knowledge” (P01).

Although they were not mentioned as much in the daily reflections, teachers highly valued the ethics activities as powerful opportunities for reflection.

“Having them be able to see on paper the different people involved, that’s going to be valuable...It helps them ask questions about why things are made and why they’re the way they are.” (P02)
 “It [was] the perfect mix: ‘Here’s code. Here’s the ethics. How could you apply it? [What are the] positive and negative effects?’” (P03).




Figure 4: Each day, we asked our 7 teachers which activities were most difficult for them or their students. The x-axis is a count of how many teachers mentioned a specific moment and the y-axis refers to the day a specific moment occurred.

Effectiveness of the Technology

Teachers felt that the Scratch software and micro:bit robots were both important parts of students learning about AI. The additional functionality provided by our custom Scratch blocks and machine learning extensions was seen as a great improvement to official Scratch. It was beneficial to use Scratch since many students were familiar with it.

“Customized Scratch allowed me to [understand AI] and introduce it to 5th graders.” (P06)
 “Making it Chromebook accessible was perfect for this time” “Scratch has been in the school for a long time.” (P07)

The micro:bit robots were viewed as a perfect complement to Scratch - providing a hook for students and enabling opportunities for collaboration.

“This is a great opportunity to be like ‘you write the code’ and ‘you do the physical part’” (P05)
 “My autistic student latched on to the robotics.” (P06)
 “[I] liked the integration of physical computing with the micro:bit and robot.” (P07)

Supporting Students

There was tension between making the curriculum compact enough to fit into a public school class period and making it accessible to students at different levels. Teachers navigated these tensions by spending extra time with students after class and even including parents.

“Some students were a bit younger [and] new to Scratch...I had to show them the pieces and how to use them. Once they got the general idea they were playing around with it.” (P03)
 “What worked for me, we went on break and I went back and retaught...Their confidence was built up by getting it once they were taught the way they needed to be taught.” (P06)

Even though teachers had these strategies, their experiences highlight the importance of making curricula as flexible as possible, so teachers can adjust them to meet their students’

needs.

The online and project-based nature of this course made it difficult to support students when they got stuck.

“They would respond that they were fine only to find out that they didn’t save the program or they had lingering questions.” (P01) “I felt like I was pulling teeth trying to get them to talk.” (P04)

In some cases, as we mentioned earlier, students were just demonstrating persistence. However, others seemed to be struggling with wanting to fit in and keep up. To overcome these challenges, teachers emphasized the need for relationship building before and coaching students on how to ask for help.

Finally, we saw teachers supporting students in developing a STEM identity, an important yet often unnamed attitude to develop in AI education (Lao 2020). Many teachers elected to participate in this curriculum because they work with underrepresented students in tech. They supported students by reinforcing growth mindset and encouraging students to explore their interests.

“I want students to feel that they have a place in STEM even if they choose other fields.” “[Being successful] doesn’t take perfection. It doesn’t mean that they don’t belong if they don’t know as much as the next person” (P07)

Teachers noted that the teaching staff being primarily women and the two researchers both being women of color positively impacted students’ perceptions of their belonging.

Teachers’ Impressions of the Curriculum

Given the overall success of the curriculum, teachers were excited to bring *How to Train Your Robot* to their schools. In Fall 2020, six teachers reported using some of the activities with their students. Vazhayil et al. also observed teachers’ excitement to begin using AI curricula right away after they had been trained to use it (Vazhayil et al. 2019).

“I’ve had experience teaching an overview of AI with high school kids but I’ve never discussed it with middle schoolers. With this I think I could.” (P03)

“Planning on incorporating this into what I already do - all of it...it’s very informational. For all of the kids to use this stuff is very important” (P04).

The biggest change they would make is extending the curriculum so students, especially new programmers, could better digest the new information. Teachers reported that when students struggled it was mostly with programming (see Figure 4). They suggested a prerequisite course in Scratch for new programmers.

“More time with the Scratch activities especially if they’ve never used it before. So they can learn how to use it better before they get into the AI stuff.” (P04)

“My biggest thing is I would stretch this out. There was more here but a week is not enough time.” (P05)

“Maybe a week before class sending students key terms and terminology homework so they can get familiar.” (P06)

Discussion

Our experience designing this course and receiving feedback from real teachers who tried the course provided us with major insights on how to create activities, design tools, and support students. In the future, we would like to explore additional tools teachers suggested like using pre-assessments, to better measure students’ learning, and finding ways to connect ideas with art and future career possibilities.

Availability concerns involving students’ economic background, geographic location, and cultural background also came up. We required students to have a laptop, preferably a Chromebook. But, future work should consider how older computers, community computers, and mobile devices could be used to lower the technology barrier. Relatedly, the digital divide, especially in rural areas, is an unsolved issue that makes computer science education unattainable for many (Reisdorf et al. 2019). Our platform uses very little Internet bandwidth. Still, some teachers and students struggled due to unreliable Internet connections. Researchers need to develop more AI platforms for schools that do not have or do not allow access to the Internet (Vazhayil et al. 2019).

The cost of robots was also a concern. Teachers said that they would either group more students per robot or do the activities without the robot parts. Although we provide a robot sprite in the programming platform to address this issue, research shows that virtual manipulatives are imperfect replacements for physical ones (Zacharia and Olympiou 2011; Moyer, Bolyard, and Spikell 2002). Future work could explore how objects already in students’ environments could help them explore AI. Finally, language accessibility was also an issue we ran into. We translated our platform from English into Spanish, but more work needs to be done to translate not just words, but the cultural metaphors that we use to make connections with students. Future research could explore co-designing AI platforms with teachers to support their students’ cultural backgrounds and learning needs.

Conclusion

In the future, we look forward to more K-12 AI curricula that strive to empower the most under-served students in the most hard-to-reach spaces. We designed the *How to Train Your Robot* curriculum, activities, technology, training with K-12 public schools in mind. These efforts allowed us to reach classrooms and build with educators and experts who understand students’ needs. We provided teachers with the curricula, materials, and training they needed to run AI workshops themselves. They successfully led their students through hands-on, technical, and ethical activities to equip them with the tools students need to be able to both build and critique AI. As researchers create curricula covering different AI topics, future interventions should address accessibility concerns like geographical location, cultural compatibility, and cost.

References

Bhatia, N.; and Lao, N. 2020. Using Transfer Learning, Spectrogram Audio Classification, and MIT App Inventor

to Facilitate Machine Learning Understanding. *CoolThink@JC* 108.

Carney, M.; Webster, B.; Alvarado, I.; Phillips, K.; Howell, N.; Griffith, J.; Jongejan, J.; Pitaru, A.; and Chen, A. 2020. Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–8.

De La Higuera, C. 2019. A report about education, training teachers and learning artificial intelligence: Overview of key issues. http://cdlh7.free.fr/UNESCO/Teaching_AI-cdlh.pdf. “Not published”.

DiPaola, D.; Payne, B. H.; and Breazeal, C. 2020. Decoding design agendas: an ethical design activity for middle school students. In *Proceedings of the Interaction Design and Children Conference*, 1–10.

Druga, S. 2018. *Growing up with AI: Cognimates: from coding to teaching machines*. Master’s thesis, Massachusetts Institute of Technology.

Fiesler, C.; Garrett, N.; and Beard, N. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, 289–295.

Heinze, C.; Haase, J.; and Higgins, H. 2010. An action research report from a multiyear approach to teaching artificial intelligence at the k6 level. In *1st symposium on educational advances in artificial intelligence*.

Hitron, T.; Orlev, Y.; Wald, I.; Shamir, A.; Erel, H.; and Zuckerman, O. 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.

Hitron, T.; Wald, I.; Erel, H.; and Zuckerman, O. 2018. Introducing children to machine learning concepts through hands-on experience. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*, 563–568.

Jordan, B.; Devasia, N.; Hong, J.; and Williams, R. 2021. PoseBlocks: A toolkit for creating and dancing with AI. In *The 11th Symposium on Educational Advances in Artificial Intelligence*.

Kahn, K.; Megasari, R.; Piantari, E.; and Junaeti, E. 2018. AI Programming by Children using Snap! Block Programming in a Developing Country. In *EC-TEL Practitioner Proceedings 2018: 13th European Conference On Technology Enhanced Learning, Leeds, UK, September 3-6, 2018*. URL <http://ceur-ws.org/Vol-2193/paper1.pdf>.

Koski, M.-I.; Kurhila, J.; and Pasanen, T. A. 2008. Why using robots to teach computer science can be successful: theoretical reflection to andragogy and minimalism. In *Proceedings of the 8th International Conference on Computing Education Research*, 32–40. ACM.

Kumar, A. N. 2004. Three years of using robots in an artificial intelligence course: lessons learned. *Journal on Educational Resources in Computing (JERIC)* 4(3): 2.

Lane, D. 2021. *Machine Learning for Kids: A Project-Based Introduction to Artificial Intelligence*. San Francisco, CA, USA: No Starch Press. ISBN 1718500564. “In press”.

Lao, N. 2020. *Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens*. Ph.D. thesis, Ph.D. Dissertation. Massachusetts Institute of Technology.

Long, D.; and Magerko, B. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16.

Marques, L. S.; Gresse von Wangenheim, C.; and Hauck, J. C. 2020. Teaching Machine Learning in School: A Systematic Mapping of the State of the Art. *Informatics in Education* 19(2): 283–321.

McConnell, J. J. 1996. Active learning and its use in computer science. In *Proceedings of the 1st conference on Integrating technology into computer science education*, 52–54.

McNally, M. F.; and Klassner, F. 2007. Demonstrating the Capabilities of MindStorms NXT for the AI Curriculum. In *AAAI Spring Symposium: Semantic Scientific Knowledge Integration*, 103–104.

Moyer, P. S.; Bolyard, J. J.; and Spikell, M. A. 2002. What are virtual manipulatives? *Teaching children mathematics* 8(6): 372–377.

National Council of Supervisors of Mathematics. 2013. Improving Student Achievement in Mathematics by Using Manipulatives with Classroom Instruction. *Position Papers: Improving Student Achievement*.

Papert, S. 1980. *Mindstorms: Children, computers, and powerful ideas*. Basic Books, Inc.

Payne, B. H. 2019. An Ethics of Artificial Intelligence Curriculum for Middle School Students. <https://www.media.mit.edu/projects/ai-ethics-for-middle-school/overview/>.

Payne, B. H. 2020. *Can my algorithm be my opinion?* Master’s thesis, MIT.

Reddy, T.; Williams, R.; and Breazeal, C. 2021. Text classification for AI education. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education, SIGCSE ’21*. New York, NY, USA: Association for Computing Machinery.

Reisdorf, B. C.; Yankelevich, A.; Shapiro, M.; and Dutton, W. H. 2019. Wirelessly bridging the homework gap: Technical options and social challenges in getting broadband to disconnected students. *Education and Information Technologies* 24(6): 3803–3821.

Resnick, M.; Maloney, J.; Monroy-Hernández, A.; Rusk, N.; Eastmond, E.; Brennan, K.; Millner, A.; Rosenbaum, E.; Silver, J.; Silverman, B.; et al. 2009. Scratch: programming for all. *Communications of the ACM* 52(11): 60–67.

Rose, D. H.; and Meyer, A. 2006. *A practical reader in universal design for learning*. ERIC.

Talaga, P.; and Oh, J. C. 2009. Combining AIMA and LEGO mindstorms in an artificial intelligence coursetobuild real-worldrobots. *Journal of Computing Sciences in Colleges* 24(3): 56–64.

Tang, D.; et al. 2019. *Empowering novices to understand and use machine learning with personalized image classification models, intuitive analysis tools, and MIT App Inventor*. Master's thesis, Massachusetts Institute of Technology.

TechGirlz. 2018. A Look Inside The Human Neural Network Activity From Our 'Artificial Intelligence: How Computers Learn' TechShop. <https://www.techgirlz.org/a-look-inside-the-human-neural-network-activity-from-our-artificial-intelligence-how-computers-learn-techshop/>.

Touretzky, D.; Gardner-McCune, C.; Martin, F.; and See-horn, D. 2019. Envisioning AI for K-12: What should every child know about AI? In *AAAI*.

Touretzky, D. S. 2014. Teaching Kodu with physical manipulatives. *ACM Inroads* 5(4): 44–51.

Tuomi, I. 2018. The impact of artificial intelligence on learning, teaching, and education. *Luxembourg: Publications Office of the European Union*.

Vahrenhold, J.; Cutts, Q.; and Falkner, K. 2019. Schools (K-12). *The Cambridge Handbook of Computing Education Research*, Sally A. Fincher and Anthony V. Robins (Eds.). Cambridge University Press, Cambridge 547–583.

Vazhayil, A.; Shetty, R.; Bhavani, R. R.; and Akshay, N. 2019. Focusing on Teacher Education to Introduce AI in Schools: Perspectives and Illustrative Findings. In *2019 IEEE Tenth International Conference on Technology for Education (T4E)*, 71–77. IEEE.

Williams, R. 2021. How to Train Your Robot: Project-Based AI and Ethics Education for Middle School Classrooms. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education, SIGCSE '21*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380621. doi:10.1145/3408877.3439690. URL <https://doi.org/10.1145/3408877.3439690>.

Zacharia, Z. C.; and Olympiou, G. 2011. Physical versus virtual manipulative experimentation in physics learning. *Learning and instruction* 21(3): 317–331.

Zimmerman, M. 2018. *Teaching AI: Exploring New Frontiers for Learning*. International Society for Technology in Education.

Zimmermann-Niefield, A.; Turner, M.; Murphy, B.; Kane, S. K.; and Shapiro, R. B. 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.