Impact.AI

Democratizing AI through K-12 Artificial Intelligence Education

by

Randi Williams

B.S., University of Maryland, Baltimore County (2016)
S.M., Massachusetts Institute of Technology (2018)

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the
Massachusetts Institute of Technology
February 2024

Authored by: Randi Williams
Program in Media Arts and Sciences
January 17, 2024

Certified by: Cynthia Breazeal
Professor, Program in Media Arts & Sciences
Dean of Digital Learning

Accepted by: Joseph Paradiso
Academic Head, Program in Media Arts and Sciences

Impact.AI
Democratizing AI through K-12 Artificial Intelligence Education

by

Randi Williams

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, on January 17, 2024 in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Abstract

Today's youth are growing up in a world where artificial intelligence (AI) technologies shape how we live, work, play, socialize, and navigate our world. This rapid technological change is already significantly shifting individuals' lives and the opportunities they can obtain. Thus, researchers, educators, and government leaders must consider how to prepare a diverse citizenry to thrive in the emerging age of AI, for example, through outreach initiatives like grade school AI curricula. My thesis delves into K-12 AI literacy, particularly how AI curricula might empower students to see themselves as technosocial change agents, capable of using technology to work toward positive, equitable social change.

First, I explore the question,"What should K-12 youth know about AI?" and introduce a new AI literacy framework, Impact.AI, covering AI concepts, practices, and perspectives that align with a technosocial change agent identity. This framework will inform the development of middle school AI curricula that empower students to become conscious consumers, ethical engineers, and informed advocates of AI. Next, I consider, "How should we design AI curricula for K-12 students and educators?" and share how I iteratively developed AI education tools and curricula that facilitate students' learning about AI as they work on AI projects. Finally, I evaluate how well these frameworks and designed artifacts contributed to students' learning about AI and developing strong AI identities. As AI becomes increasingly prevalent in everyday life, it is essential that all people have the opportunity to both understand and shape the technology.

Thesis Advisor:
Cynthia Breazeal

Impact.AI

Democratizing AI through K-12 Artificial Intelligence Education

by

Randi Williams

This dissertation has been reviewed and approved by the following committee members

Advisor: ______________________________

Dr. Cynthia Breazeal
Professor of Media Arts and Sciences, Dean of Digital Learning
Massachusetts Institute of Technology

Committee Member : ______________________________

Dr. Hal Abelson
Class of 1922 Professor of Computer Science and Engineering
Massachusetts Institute of Technology

Committee Member : ______________________________

Dr. Tia C. Madkins
Assistant Professor of Curriculum and Teaching
The University of Texas at Austin

Committee Member : ______________________________

Dr. Jean Ryoo
Director of Research, Computer Science Equity Project, Center X
University of California, Los Angeles

Contents

List of Figures

List of Tables

Acknowledgments

Hold fast to dreams
For if dreams die
Life is a broken-winged bird
That cannot fly.

Hold fast to dreams
For when dreams go
Life is a barren field
Frozen with snow.

Langston Hughes

This dissertation is dedicated to dreams.

To Dr. Cynthia Breazeal

I must profess the most profound gratitude to you, who made this work possible on multiple levels - from your vision to financial support, connections, guidance, critique, comfort, and unwavering support. You made this work possible. You inspired me and helped me achieve a beautiful dream. My graduate school experience was so fulfilling; I know how lucky that makes me. Thank you for being a fantastic advisor.

To Dr. Hal Abelson

Thank you for years of encouragement and mentorship. Your excitement to empower all people with computational action is contagious. I am so proud of how the work we have done through RAISE has impacted grown up and impacted lives. You are indeed a hero in the field; thank you for your leadership.

To Dr. Tia Madkins and Dr. Jean Ryoo

Thank you for choosing to be on my team for this dissertation. I have been a fan from afar for years, watching how you reimagined what CS could be for everyone, especially the people who keep getting left out. The opportunity to meet with you and have you shape and stretch my ideas was my honor. Thanks to both of you, my skills and belief in myself as a researcher have evolved on multiple levels.

To the Media Lab

...and especially to the Personal Robots Group, thank you for being my homies. Thank you for countless days and nights of dreaming and scheming to make the world a better place. My fellow students have been incredible collaborators and partners in conducting studies together, building out our wacky ideas, spending long nights coding or writing papers, and savoring every joyful moment along the way. I am particularly grateful to the AI education team, whom I thoroughly cited throughout this dissertation.

Without its incredible facilities, administration, member relations, and communications staff, the Media Lab would be nothing. Thank you for welcoming, supporting, pushing, surviving graduate school during a pandemic, and teaching me to see the world in infinite new dimensions.

To my family and friends

Parents, thank you for first sowing the seed that grew into this dissertation. Sisters, thank you for helping nurture it. Because of your examples and encouragement, I saw myself as a technosocial change agent at a very young age. It is an honor to pass on your legacy. To all my family and friends, thank you for hearing me, cheering me, and feeding me so I could keep going. I appreciate you all more than I could ever express.

To Prince George's County, the Meyerhoffs, Baltimore, and D.C.

This dissertation and all my years of graduate school are dedicated to the communities that shaped me. Thank you for teaching your children to be strong, to reach back for others, and to dream big dreams. You are the reason why I do what I do.

1 Introduction

We live in a world where artificial intelligence (AI) technologies have permeated almost every aspect of our daily existence, and their influence is only growing. It is unimaginable to fathom that smart speakers only emerged in 2012 and then, ten years later, were regularly used by nearly half of the United States population [1]. Then, in November 2022, OpenAI made ChatGPT[1] available to the public and it broke records as the fastest-growing web app ever, reaching more than 1 million users in just five days [2].

Data-driven algorithms increasingly influence many aspects of our lives – from social interactions to financial opportunities, from entertainment to healthcare. This has created enormous opportunities for us to explore, enhance, and reimagine how we navigate our world. These record-breaking usages and incredible opportunities come with their share of challenges.

This introduction outlines three major concerns linked to the emergence of AI systems: the digital literacy gap, the Computer Science Education (CSED) gap, and the AI ethics crisis. Understanding these challenges provides context for the contributions made in this work and sets the tone for the remaining chapters of this dissertation.

1: ChatGPT, https://chat.openai.com/chat

1.1 The Digital Literacy Gap

Understanding children's changing relationships with AI

The presence of AI in everyday life is having a profound impact on today's youth. Early during my time at the Media Lab, I supported a project that used a social robot to enhance preschool-aged children's reading and writing skills (Figure 1.1). The literacy intervention involved students interacting with the social robot to practice reading, pronouncing words, and expanding their vocabulary.

Figure 1.1: The author, Randi Williams, sitting before Tega the social robot, set up for a language literacy intervention. Learn more about Tega at https://www.youtube.com/watch?v=U4srV1Icnb0

One interaction with a student went like this:

> Randi (author), starting the intervention: Hi, my name is Randi. Today, you will be interacting with a social robot named Tega. Let me show you how Tega works.
> Student, to Randi: I already know how to speak with robots. I have an Alexa at home.
> Student, to robot: Hey, Tega. What is the weather like today...
>
> *Paraphrased from a conversation with a Cambridge public school student*

This encounter left a lasting impression on me due to its implications for the evolving relationships between children and technology. I reflected on my experiences learning to navigate the Internet by"asking Jeeves" various questions. Back then, searching the Internet required someone to be able to read, write, type, log on to the Internet, and then phrase questions in a manner that Jeeves could understand. In stark contrast, today's children can use speech interfaces to explore the Internet before learning to read or write.

Comparing the experience of asking Jeeves a question with querying Alexa highlights a substantial shift in how we relate to technology. The natural developmental tendency to anthropomorphize is more potent when devices are more, well, anthropomorphic, closely mirroring the ways that humans communicate with one another [3–6]. This means that Alexa, a named entity with a speech interface, feels more lifelike and human-like than Jeeve's web page. Additionally, there is the matter of how knowledge is conveyed: while asking Jeeves yields thousands of websites that one could browse to synthesize an answer, Alexa provides a single response. The answers provided by Alexa, whether facts or opinions are delivered with an air of confidence and finality, making the device seem more intelligent and the answers more trustworthy. With the emergence of new intelligent technologies (I'm looking at you, ChatGPT), there are questions about how these tools shape young people's relationship with technology and, consequently, the world will continually evolve.

As AI systems become increasingly pervasive, persuasive, and alluring, consumers must acquire the skills they need to use them safely and effectively. Users must be able to safeguard their privacy, discern credible sources from misinformation, and maintain their agency. Yet, a Gallup and Northeastern University survey revealed that most adults lack basic AI knowledge [7]. These concerns intensify for younger users due to research showing their tendency to form trusting relationships with personified agents [8]), putting them at risk for privacy and consumer rights violations [9].

Despite growing awareness of the risks of AI among consumers, as indicated by a Common Sense Media survey where 72% of teens stated they believed that social media and video hosting site advertisements were intentionally designed to manipulate them [10], the lack of transparency in AI algorithms remains a significant concern. With limited age-appropriate opportunities for youth to learn about AI, educators, and researchers are calling for expanding AI literacy initiatives for all ages.

In my Master's thesis, I delved into children's perceptions of intelligent agents and then subsequently designed PopBots, programmable social robots created to help preschool children learn about AI [11]. Experimental studies with PopBots revealed that the preschool AI lessons deepened participants' deeper understanding of intelligent agents, encouraging them to think more critically about examples of AI in their lives. However, their perspectives about AI were also shaped by factors including their prior experiences with technology, media depictions of intelligent machines, and how others discussed the technology. Building on these findings, my dissertation helps bridge the digital literacy gap by utilizing youths' understanding of technology, the world, and themselves as a springboard, preparing them to become technically adept and socially conscious creators of AI.

[11]: Williams (2018), 'PopBots: leveraging social robots to aid preschool children's artificial intelligence education'

1.2 The Computer Science Education Inclusion Gap

Who is in the pipeline, and who is left out?

While education offers a promising avenue to empower children to transition from AI consumers to creators, there are significant challenges to be addressed. Ensuring that all people have equal opportunities to benefit from AI requires concerted efforts to overcome barriers to access and participation. Initiatives to expand Computer Science Education (CSED) in K-12 schools have been ongoing for over two decades, with some progress. According to the 2022 Code.org State of CS Report, today, over half of high schools in the United States offer foundational computer science classes to their students [12]. However, there is still work to be done to ensure equitable access for all.

[12]: Code.org et al. (2022), '2022 State of Computer Science Education: Understanding Our National Imperative'

Women, racial minorities, and other students from marginalized backgrounds still are not participating in K-12 CS courses at the same rate as their peers. The reasons for this under-participation are complex and persistent. For one, these students have to contend with stereotypes about who can be successful and STEM that develop early in childhood, long before high school [13, 14]. Unfortunately, access to CSED in primary and middle schools, crucial periods when students form beliefs about their potential in CS, is unequally distributed. On top of this, students must wrestle against stereotype threat, micro-aggression, and other forms of everyday discrimination that undermine expectations for success and feelings of belonging [15].

Another part of the issue is structural. [16]. As highlighted by the COVID-19 pandemic that began in 2020, there are huge disparities in funding, school resources, and access to technology that negatively impact the educational outcomes of many students [17]. The result is that students who are racialized as Black/African-American, Hispanic/Latine/Latina/Latino/Latinx, Native American/Alaskan and American Indian, Native Hawaiian/Pacific Islander, are less likely to go to schools that offer CS courses. On top of this, a students' race as well as gender could also mean they are statistically less likely to participate in a CS course when it is available. Similar challenges persist for multilingual English language learners, students with cognitive and physical disabilities, and economically disadvantaged students [12].

To overcome these barriers, there are many potential paths forward. A structural path is policy. Code.org has identified 9 policies for making CS, and states that have adopted at least 6 of these 9 policies have much higher proportions of schools that offer CS. These policies can also address participation gap issues, as the three states with the highest numbers of women participating in computer science courses have policies that make CS a graduation requirement or an option to satisfy a requirement. Another direction for progress involves addressing the barriers that keep students from having access to and participating in these courses. The Code.org report identifies a lack of funding and access to technology, a lack of appropriate curricula, and a lack of teacher training as important barriers to address [12]. A study of CS teachers done by the Kapor Center highlights that equitable access also means promoting CS education that is culturally relevant and empowering [16].

[16]: Kapor Center (2021), 'Culturally responsive-sustaining computer science education: A framework'

AI represents a new frontier of CSED, allowing us to learn from past errors and improve our approach. Code.org pinpointed three key obstacles hindering access to CS education: the absence of suitable curricula, insufficient funding and technological resources, and inadequate teacher training. This dissertation addresses these challenges and centers the experiences of students from communities that have been historically marginalized by the tech industry. Co-designed with educators and students, the free, open-source curricula and tools that precipitated from this dissertation prioritize students' interests and are specifically designed to be compatible with their classroom environments. Through educator training programs, the curricular materials were refined to ensure effective engagement. My objective was to establish a model for achieving equitable AI education, open for adaptation and critique, to enable broader access to AI education for diverse communities.

1.3 The AI Ethics Crisis

AI for who's good?

In a world increasingly shaped by AI, fostering inclusive participation in the creation and critique of these systems is paramount. Despite being portrayed as neutral and progressive, these technologies often recreate and exacerbate systems of harm that we as a society are trying to move away from [18, 19]. Part of this struggle is that the lack of diversity in the AI workforce persists as a weakness of the field. As shown in Table 1.1, women and minoritized racial groups are extremely underrepresented in AI positions [20]. Data on representation by educational background, sexual orientation, gender minority status, and disability status are less forthcoming, but the pattern of exclusion likely extends to these identities as well.

[20]: West et al. (2019), 'Discriminating systems'

Table 1.1: Data collected by [20] on gender, race, and participation in the AI workforce.

Academia
Female professors in AI - 20%
Female authors in top AI conference proceedings - 18%
Industry
Female research staff at:
Facebook - 15%
Google - 10%
Black/ African - American full-time employees at:
Facebook - 4%
Google - 2.5%
Microsoft - 4%
Hispanic/Latine/Latinx full-time employees at:
Facebook - 5%
Google - 3.6%
Microsoft - 6%

Organizations like OpenAI and Diversity.AI have raised concerns about the concentration of AI knowledge in the hands of a few who do not represent the full diversity of all impacted parties. Researchers have shown how limited diversity in the design and evaluation of AI systems sometimes directly leads to deployed systems that illegally discriminate based on demographic factors like race, sex, income, and age [21–23]. Addressing this issue requires reevaluating the AI pipeline: who is represented and who is not? Which values, cultures, norms, and perspectives dominate the field, and which ones have been continually overlooked?

> "There is a bias to what kinds of problems we think are important, what kinds of research we think are important, and where we think AI should go. If we don't have diversity in our set of researchers, we are not going to address problems that are faced by the majority of people in the world."
>
> Timnit Gebru
> *"We're in a diversity crisis": co-founder of Black in AI on what's poisoning algorithms in our lives*, 2018

Beyond questions around representation, it is also vital to address challenges uniquely faced by marginalized individuals who participate in the

AI field. Turner, Wood, D'Ignazio, and Teng (2021) highlighted alarming patterns where institutions shame, silence, and erase Black women whose contributions to AI challenge dominant narratives, particularly when they call out examples of inequity in the field [24]. Black women, women from low-income backgrounds, people within the Disability community, and LGBTQ folks are often at the forefront of AI ethics work, revealing algorithmic harms caused by deployed systems and creating action to develop more inclusive systems that truly benefit all people. However, this work can come at significant personal cost to those who speak out unless individuals and the field challenge the values and structures that perpetuate social inequality.

[24]: Turner et al. (2021), 'The abuse and misogynoir playbook'

A common view held amongst those in the tech industry is that technology can be used to solve problems and make the world a better place. Movements like AI for Social Good (AI4SG) leverage AI technology to work toward social equality, improving quality of life, and expand access to resources. However, even well-intentioned AI systems can backfire if systems designers fail to attend to ways that societal bias may be built into every step of the AI development process [25]. To ensure that tech reaches its goals of improving society, rather than exacerbating social harms, AI practitioners must prioritize inclusion and ethics in their work.

This dissertation addresses the AI ethics crisis by drawing on scholarship from and about marginalized communities that seek to transform AI's design process. Incorporating educational frameworks like Culturally Responsive Computing [26] and Universal Design for Learning [27], alongside human-centered design philosophies like design justice [28], this dissertation builds on rigorous practices for achieving equity. By proactively addressing AI's potential harms, this dissertation forges new paths toward social progress where everyone's well-being is fundamental to technological success.

[26]: Scott et al. (2015), 'Culturally responsive computing: A theory revisited'

[27]: CAST (2018), 'Universal Design for Learning Guidelines version 2.2'

[28]: Costanza-Chock (2020), *Design justice: Community-led practices to build the worlds we need*

1.4 Contributions of This Dissertation

> "I celebrate teaching that enables transgressions - a movement against and beyond boundaries. It is that movement which makes education the practice of freedom"
>
> bell hooks *Teaching to transgress*, 1994

This dissertation employs education as a tool to tackle the major challenges in digital literacy, participation in CSED, and attention to AI ethics that are currently limiting AI's progress. In partnership with middle school educators, I iteratively designed AI curricula that equip students with the knowledge and perspectives that correspond to a technosocial change agent identity. Through evaluative studies, I demonstrated that these curricula effectively enhanced students' technical and ethical understanding of AI, equipped them with the skills to create AI projects with real-world impact, and positively impacted their attitudes toward AI. Concurrently, I developed AI-powered education tools that aid educators in nurturing their students' technical proficiency, critical thinking, and creativity as they learn about AI and apply their knowledge to real-world projects.

I specifically focus on AI + ethics curricula for middle school students because, as pointed out by Blakeley H. Payne (2020) in her thesis about the first middle school AI + ethics curriculum, middle school students are developmentally prepared to navigate ethical dilemmas and independently develop projects for real-world impact [29]. This work deliberately centers students from backgrounds that the tech field has historically marginalized and harmed, students who in the literature are often referred to as “excluded" and “excoded" [30]. Contrasting the current status quo of race-neutral AI curricula, that can ultimately perpetuate harm and exclusion by overlooking the needs of the most vulnerable students [31, 32], this work “includes" and “encodes" all students toward a more equitable AI future.

[29]: Payne (2020), ‘Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students’

[30]: Buolamwini (2022), ‘Facing the Coded Gaze with Evocative Audits and Algorithmic Audits’

Contributions to AI Literacy The contributions to AI Literacy include the design of the first culturally relevant AI literacy framework (Chapter 3), plus multiple culturally relevant AI curricula (Chapters 4 & 5). The curricula are free, publicly available, and open to use and adaptation in non-commercial ventures.

Contributions to AI Education The contributions to AI education include the development of a hands-on AI education resources, including a block-based AI programming platform and programmable robots (Chapter 4), plus integrated scaffolding tools that support students’ technical skills, ethical reasoning, and creativity (Chapter 6). The AI programming and scaffolding tools are free, open source, and work on any Internet-connected device with a browser.

Contributions to AI for Social Good The contributions to AI for social good include a series of AI education artifacts (Chapters 4 & 6) and collection of AI projects (Chapters 5 & 6) developed with, by, and for middle school students and educators. The participants in this work are primarily from groups that continue to be underrepresented in and excluded from tech, thus the work presented throughout this dissertation offers unique perspectives on AI development to can enrich the field.

2 Background and Related Work

This dissertation draws upon three bodies of literature: K-12 AI literacy, diversity and inclusion in Computer Science Education (CSED), and AI's applications in education. In this section, we:

- Examine existing frameworks on K-12 AI literacy, exploring their motivations and recommendations on student learning objectives
- Overview published literature on diversity and inclusion in education, particularly examining culturally responsive pedagogy and its implications for AI literacy
- Survey AI education technologies that support innovative pedagogical approaches

This literature review contextualizes the importance of the work accomplished in this dissertation and reveals opportunities for future work.

2.1 K-12 AI Literacy

2.1.1 K-12 AI Literacy Frameworks

In recent years, researchers have introduced various frameworks on AI literacy for learners in the general public. Three notable examples are the blue sky paper on the Five Big Ideas in AI by Touretzky et al. (2019) [33], the AI Literacy definition paper by Long and Magerko (2020) [34], and the UNESCO AI Competency Framework for Students and Teachers [35].

[33]: Touretzky et al. (2019), 'Envisioning AI for K-12: What should every child know about AI?'
[34]: Long et al. (2020), 'What is AI literacy? Competencies and design considerations'
[35]: UNESCO (2023), 'Draft AI competency frameworks for teachers and for school students'

These three frameworks articulate different motivations for educating K-12 students about AI. Touretzky et al. (2019) [33] underscore the importance of empowering students as informed consumers and potential contributors to AI, aligning their framework with the widely adopted CSTA K-12 Computer Science Standards. The Five Big Ideas in AI framework uniquely tailors AI competencies for specific grade bands K-2, 3-5, 6-8, and 9-12. Next, Long and Magerko (2020) identified 17 competencies for achieving AI literacy, informed by their review of existing K-12 AI instruction resources [34]. In addition to specifying what students should learn, their AI literacy framework also recommends design principles to guide the creation of future curricula. Finally, The UNESCO AI competency framework[1] articulates the knowledge, skills, and attitudes essential for students to safely engage with AI. UNESCO's framework is the only one that presents frameworks for both teacher and student AI literacy.

Each framework describes the knowledge they would like learners to acquire, recognizing the importance of tailoring this to learner's interests and needs. All three frameworks prescribe that students should learn a range of AI topics, covering foundational knowledge of AI, specific concepts, problem-solving methods, and context for societal impact. The

Five Big Ideas in AI framework categorizes all of AI into five areas, interspersing specific concepts and skills among these categories. The AI Literacy and UNESCO frameworks are organized around specific skills students should learn, including applying knowledge, programming, and critiquing AI systems, across subtopics within AI. The UNESCO framework uniquely identifies AI attitudes that students should develop, particularly a human-centered mindset where students critically reflect on AI, consider how to use it responsibly, and develop a sense of their power to shape the technology[1].

Beyond K-12 AI education, there are additional frameworks worth exploring that can inform what K-12 students should learn about AI. Brennan and Resnick (2012) published a taxonomy of critical computational thinking knowledge that K-12 learners should develop, called the concepts, practices, and perspectives (CPPs) framework [36]. Concepts are defined as the set of knowledge that practitioners engage with as they do their work, practices involve the skills and methods that practitioners use, and perspectives capture evolving beliefs that practitioners develop as they complete their work. The CPP framework has been used to develop measurable learning outcomes and effective assessments for computational thinking curricula. Similarly, in this dissertation, a CPP framework for AI literacy is used to organize AI learning and assessments for effective AI curricula.

[36]: Brennan et al. (2012), 'New frameworks for studying and assessing the development of computational thinking'

Another beneficial framework, that also employs a CPP-like structure, is the ML Education framework for "tinkerers and ML-engaged citizens" developed by Lao (2020) [37]. This framework summarizes machine learning education by the AI knowledge, skills, and attitudes (correlates to concepts, practices, and perspectives) that learners should develop. Lao (2020) emphasizes the importance of including students' developing attitudes about AI in their AI literacy framework. Attitudes are important because learners need to develop self-efficacy and a sense of belonging in the field to become efficacious contributors to the field [37]. The ML Education framework also attends to students gaining knowledge about AI that enables them to be optimistic critics and responsible advocates of AI. The framework has students learn about how how different kinds of bias propagate through the entire machine learning pipeline and then proactive methods for reducing the harm of these biases [25, 38]

[37]: Lao (2020), 'Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens'

[37]: Lao (2020), 'Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens'

The Impact.AI framework detailed in Chapter 3 closely resembles the structures of the UNESCO AI student framework, CPP framework, and ML Education framework. Impact.AI encompasses the concepts, practices, and perspectives that are essential for K-12 students to comprehend and develop. While sharing motivations with these frameworks, Impact.AI uniquely integrates Culturally Responsive-Sustaining Pedagogies (CRPs) to empower students in recognizing their ability to use AI to transform society. The other frameworks express the importance of inclusivity and broadening participation in AI, but none explicitly leverages pertinent social or educational theories for educational equity. Impact.AI seeks to address this gap by treating students' identities as highly relevant, rather than incidental, to the instruction they need [31]. The next background section delves into the imperative of supporting marginalized students to bring their complete identities into the tech field. This approach prioritizes fostering an inclusive environment where individuals do not have to conform to prevailing tech culture norms that may conflict with their authentic identities [39].

2.1.2 Trends in K-12 AI Literacy

In this section, I will summarize other literature reviews on K-12 AI literacy. This section lays the groundwork for Chapter 3, which summarizes the literature review results on K-12 AI literacy work that kicked off the development of the Impact.AI framework.

Ng et al. (2021) [40], Ojeda et al. (2021) [41], Su et al. (2022) [42], and Rizvi et al. (2023) [43] have conducted literature reviews summarizing the content and pedagogical approaches used to teach grade school students about AI. Ng et al. (2021) categorized students' knowledge into three categories: the ability to use and apply AI, skills to evaluate and create AI, and motivation to pursue a future career in AI [40] . The different pedagogical approaches they identified included discovery, projects, hands-on, and unplugged learning. Ojeda et al. (2021) cataloged AI conceptual knowledge by subject area: intelligent systems, data science, expert systems, computer vision, deep learning, and representation [41]. Su et al. (2022) used a knowledge-skills framework, distinguishing curricula based on whether they tried to convey knowledge (e.g., how AI works) or skills [42]. Finally, Rizvi et al. (2023) categorized content by whether they covered societal impact and ethical implications, applications, models, or AI engines [43].

All literature reviews pointed out a big weakness of the field: a lack of alignment on what to teach K-12 students and how to assess their knowledge [40, 42, 43]. The lack of consensus on what to teach students and how to gauge success impedes progress in determining the most effective approaches to equip students with knowledge of AI. Another concern, pointed out by reviews focused on diversity and inclusion, was that many curricula are deployed as short, extracurricular workshops [43] developed in regions with more access to the AI industry [44]. The limited reach of these workshops curtails the potential of these interventions to have a sustaining impact on broadening participation in AI.

Another series of reviews on AI curricula focused on educational tools. Ng et al. (2021) and Sanusi et al. (2021) described educational tools by whether the platform was hardware, software, intelligent agent, robot, or unplugged [40, 44]. These reviews found that most platforms support students learning about machine learning classification, computer vision, and conversational agent design, neglecting other AI topics. Gresse et al. (2021) further categorized educational tools by their pedagogical approach and design choices [45]. Many tools leveraged the Use-Modify-Create learning progression [46], which involves progressively overturning users' control as their knowledge grows. Most tools were designed for novice users without prior programming experience. Thus, tools were highly interactive and engaging, utilizing input from users' webcams, microphones, voice assistants, computer files, or other devices. Fleger et al. (2023) noted that model creation platforms often obscured many details of the machine learning process [47]. For example, one supervised machine learning tool expected users to collect data, train models, and evaluate the resultant model. However, the tool neither explained nor allowed users to tinker with the model's parameters. These reviews generally reveal a common emphasis on providing students with rich AI construction experiences while obscuring some details about how the algorithms function to make the user experience more approachable.

Our literature review highlights the extensive work in K-12 AI literacy, particularly curricula that teach machine learning to middle school aged students in informal education spaces. However, as emphasized in several reviews, there is a pressing need for more effort to evaluate curricula and identify optimal teaching approaches [40]. Until researchers begin systematically addressing reliability and validity at scale, it will be difficult for the field to gain the cohesion necessary to teach all students AI.

2.1.3 Educational Theories for K-12 AI Instruction

Like many other K-12 AI education curricula, the curricula presented in this dissertation are based on student-centered educational philosophies that prioritize hands-on learning. Impact.AI curricula incorporate constructionism [48, 49] and computational action [50], active learning [51, 52], and Universal Design for Learning [27, 53, 54]. This section describes those educational philosophies and how they are utilized in AI curricula.

[48]: Papert (2020), *Mindstorms: Children, computers, and powerful ideas*
[49]: Resnick (2017), *Lifelong kindergarten: Cultivating creativity through projects, passion, peers, and play*
[50]: Tissenbaum et al. (2019), 'From computational thinking to computational action'
[51]: Michael et al. (2003), *Active learning in secondary and college science classrooms: A working model for helping the learner to learn*
[52]: Bonwell et al. (1991), *Active learning: Creating excitement in the classroom. 1991 ASHE-ERIC higher education reports.*
[27]: CAST (2018), 'Universal Design for Learning Guidelines version 2.2'
[53]: Kieran et al. (2019), 'Connecting universal design for learning with culturally responsive teaching'
[54]: Rose et al. (2006), *A practical reader in universal design for learning.*

Constructionism and Computational Action

Constructionism is a pedagogical approach that evolved from constructivism. Constructivism posits that students learn best when actively engaged in learning activities rather than passively receiving information. Constructivism takes a student-centered approach to learning in contrast to instructivist or cognitivist theories, which emphasize how educators deliver knowledge to their students. Constructionism adds that learning can be particularly effective and engaging when students use their knowledge to create artifacts in collaboration with their peers, which they can then reflect on to further their understanding. Advocates for constructionism emphasize that students should develop personally meaningful artifacts that allow them to express and explore pieces of their identity [55]. Computational action is closely aligned with constructionism. Computational action seeks to engage students in authentic computing practice as a part of their communities [50]. This approach involves students completing significant, self-driven projects related to community concerns.

K-12 AI curricula that are aligned with constructionism and computational action empower students to create AI artifacts and complete projects that are meaningful to their communities. Pang et al. (2022) incorporated computational action into a responsible AI curriculum, working with students to define community problems, conduct user research to understand needs, and engage in design with those communities [56]. Similarly, Van Brummelen (2022) described how they used computational action in their lessons teaching high school students about conversational agents [57]. They recommend incorporating computational action into AI lessons to increase students' confidence in their ability to use AI.

Active Learning

With active learning, students play a key role in driving their learning as they engage in activities and reflect on what they have learned. Students learn as they engage in activities and then process new knowledge through reflection [51, 52]. Student agency, voice, and self-determination are prioritized through innovative pedagogical strategies like learning by design and collaborative problem-solving [52]. Reflection solidifies students' understanding of concepts by helping them clarify what they have learned and self-evaluate where they are in the learning process. Reflection practices like journaling, surveys, group discussions, and presentations give students space to articulate their knowledge.

K-12 AI curricula use various active learning techniques such as unplugged activities, simulations, and programming. AI literacy researchers have used unplugged activities to give students opportunities to explore AI algorithms even when they do not have the background mathematical knowledge or computational resources to understand the algorithm [29, 58–60]. Prior work also includes many examples of AI simulators where students and educators can interact with AI models and get a glimpse into otherwise opaque algorithms. [58, 61–67]. Finally, many existing curricula and workshops leverage novice-friendly coding tools to allow students to become designers of their own AI systems. These tools include block and text-based programming platforms that allow students to create and integrate AI components into their projects [68–71]. Given the numerous examples with developed curricula, there is much to build on for active learning in K-12 AI literacy.

Universal Design for Learning

Universal Design for Learning is an approach to education that arose from the disability justice community. It insists that educators recognize students' broad range of interests and learning styles and then provide different ways to match lessons with their strengths. Educators intentionally seek to increase students' engagement by presenting materials in multiple modalities and encouraging students to demonstrate their learning different ways. This includes considering students' prior knowledge and access to technology when designing curricula [27].

Sanusi et al. (2021) reviewed tools for teaching machine learning in K-12 [44]. They observed that most programming tools and simulations are free, web-based tools accessible to users with a stable Internet connection. This design choice is often made to include more learners by lowering the bar for access but still excludes many learners who may not have reliable Internet connections. Sanusi et al. (2021) [44] and Fleger et al. (2023) [47] also note that age-appropriateness is an important consideration in design since many tools seem focused on students in secondary school with little focus on early education. Our review of K-12 AI curricula also noted the lack of curricula focused on learners with disabilities. Bigham (2008) stood out as an example of an AI literacy curriculum that empowered blind students to learn to program conversational agents [72]. This technology could benefit people in the blind community, who are often underrepresented in creating these technologies. More work needs to be done to address access barriers, including Internet connectivity, accessibility, and opportunities to learn about AI.

2.2 Addressing Diversity in AI Literacy

As discussed in the Introduction (Chapter 1), the lack of diversity in the AI workforce is a major hurdle to the field's potential. AI ethicists have repeatedly demonstrated how the harms perpetuated by AI systems reflect broader social inequalities and the underrepresentation of vulnerable groups in the tech sector [18–20, 73]

2.2.1 Diversity and Inclusivity Initiatives in K-12 CS Ed

A common objective of K-12 AI literacy initiatives is to empower "all" individuals as users and creators of AI, with "all" being a term specifically used to invoke desires for inclusivity [34]. However, there are many barriers to achieving this goal. In CS education work, researchers recognize that the societal factors that lead to the underrepresentation of some people in the tech workforce leak into the classroom [74]. Despite the growing availability of CS courses, students from marginalized groups are underrepresented compared to their presence within the overall school population [12].

2.2.2 Culturally Responsive-Sustaining Pedagogies

Culturally responsive-sustaining pedagogies (CRPs) recognize that students need to develop academic knowledge, cultural competence, and critical consciousness to become efficacious agents for change in an inequitable society [26, 39, 75, 76]. These pedagogies fall under the broader umbrella of asset pedagogies and build on constructivism. A key idea in asset pedagogies is that all cultures core to students' identities are valuable for helping them learn [77–79]. As constructivist pedagogies, CRPs see learning as socially constructed through enculturation into a knowledge community that cannot be separated from students' prior knowledge, lived experiences, or home cultures [74].

Unfortunately, most educational systems fail to recognize that they embrace some cultures and devalue others, ultimately harming students' learning by forcing them to assimilate [80]. In culturally responsive classrooms, teachers build connections between content and students' identities, support students in their identity formation and progression, and intentionally strive for equity in their classrooms [26]. Teachers facilitate learning experiences that resonate with learners' identities, stimulating deep engagement in the topics [74, 78]. Student mastery is expressed through creation and action, and teachers give students feedback to continue and strengthen their work [81].

Frameworks such as Culturally Responsive Computing [26], Culturally Relevant CS Pedagogy [16, 39], and Liberatory Computing [82] are examples of CRPs applied to computing contexts. The Liberatory Computing Education Framework for African-American students is the only one

[26]: Scott et al. (2015), 'Culturally responsive computing: A theory revisited'

[16]: Kapor Center (2021), 'Culturally responsive-sustaining computer science education: A framework'

[39]: Madkins et al. (2019), 'Culturally relevant computer science pedagogy: From theory to practice'

[82]: Walker et al. (2022), 'Liberatory Computing Education for African American Students'

of these frameworks that focuses on an AI topic, namely data science. The framework has five pillars around identity, critical consciousness, collective responsibility, liberative identity, and activism that students engage with as they develop skills in data science methods [82]. Students completing projects in the Liberatory Computer Education applied their data-collection skills to address community problems related to disparate access to public services, exposure to natural disasters, and mortality rates. The Liberatory Computing Education framework and curriculum highlight the importance of attending to students' identity and issues of power and inequality when discussing AI topics, such that students emerged from the course with a better sense of how all of those ideas were connected. Given that all of the students in that course were African-American students, a group that continues to be underrepresented in the AI workforce and marginalized by AI, this framework has direct implications for issues of inclusion in the field.

Although there are differences in the focuses and goals of each framework, there is quite a bit of similarity in how each framework attends to identity and self-discovery, use of anti-racist practices and creation of an identity-welcoming space, attention to systemic inequality in systems, and the stressed importance of leveraging computing to affect positive change. These approaches span informal and formal educational experiences. Critical Computational Literacy is an approach to education that explicitly goes beyond traditional classrooms and into the communities and spaces where young people can affect change [76, 83]. The Culturally Relevant-Sustaining CS Education, Exploring Computer Science, and Culturally Responsive Computing frameworks all began in the classroom, focusing on training educators to build connections between CS and their students' lives and employ teaching practices that validated students identities [16, 26, 84, 85]. But many of these frameworks also move outside the classroom, incorporating students' families and communities into the learning experience to become part of their learning [16].

[76]: Lee et al. (2016), 'None but ourselves can free our minds: Critical computational literacy as a pedagogy of resistance'
[83]: Lee et al. (2023), *Code for What?: Computer Science for Storytelling and Social Justice*

This work centers the Culturally Responsive Computing framework [26], which has five tenets: (1) all students are capable of digital innovation, (2) technology should enable students to reflect and demonstrate understanding of their intersectional identities, (3) self-learning focuses on students realizing their potency as activists, (4) barometers for technological success should consider how far an individual operationalizes their technosocial agency to further their communities, and (5) in collaboration with students' communities, education should position technology as a means toward a social justice end.

In AI literacy, a culturally responsive pedagogical approach means that students learn AI competencies, develop cultural competence, attain a critical awareness of AI's role in society, and develop identities as tech change-makers who can utilize AI to create positive social change [16, 39, 86, 87]. Existing literacy frameworks usually only attend to knowledge competencies and critical awareness, and curricula have paid little attention to educational inequities that keep some groups from participating in AI. Recently, K-12 AI educators have increasingly brought culturally responsive pedagogy into their work. For example, Solyst et al. (2023) [87] and Li et al. (2023) [88], in a research collaboration between Carnegie Mellon University and the University of Pittsburgh, leveraged a culturally responsive pedagogical approach to teach students about social robotics and accountable AI.

By building on culturally responsive-sustaining pedagogies in our curricula, this work directly tackles barriers to participation that inhibit all students. Achieving equity in K-12 AI classrooms requires educators to directly address issues of access and inclusion through culturally responsive and liberatory pedagogical practices [31, 89–91]. In the next subsection, I will discuss another thread of ensuring equity in the classroom and in the field of AI by integrating ethics into AI instruction.

2.2.3 Centering Ethics and Justice in AI Instruction

Including ethics in AI instruction is extremely important and has only recently begun receiving more attention. Historically, college-level AI courses have overlooked ethics or relegated ethics instruction to the end of a course or a separate course altogether [92]. The impact of this choice is that many CS and AI engineers enter their practice without ever developing the skills to consider ethics and societal impact as they design new systems [93]. Skirpan et al. (2018) found that when students engage with ethics throughout a computer science course, they are trained to think more holistically about the implications of their work [93].

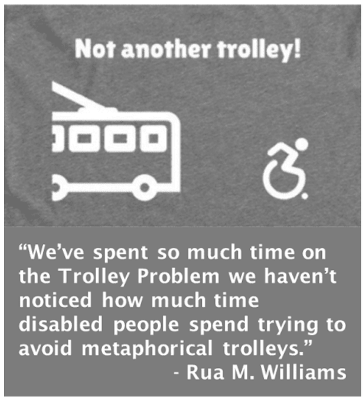

Figure 2.1: A t-shirt from the workshop in [94] featuring the words "Not another trolley!" with a quote from the paper that captures their vision for co-liberative consciousness in CS education.

Given the large, frequently harmful impact that AI has on society, and particularly on those most vulnerable, it is critical that engineers possess the skills to think about ethics. Moreover, AI developers need to be able to leverage frameworks for thinking about design and ethics as a means toward fairness and justice. Researchers like Rua Williams et al. (2021) [94] and Garrett et al. (2020) [95] pointed out that even when university computer science classes include ethics, the issues covered are often narrowly scoped to algorithmic bias in datasets rather than offering context about how bias stems from systemic patterns of social inequality. Williams et al. (2021) [94] clearly illustrate this by problematizing the "Trolley Problem" (see Figure 2.1), an ethical dilemma often posed to students in lessons about the ethics of self-driving cars [96, 97]. Rather than limiting our perspective of ethics to the fairness of making one ethical decision and deciding which life to take and which to spare, a justice-focused ethics lens would ask what patterns of harm continually put some lives in more danger than others and how technology might mitigate those systemic.

Educational movements toward ethics-centered AI literacy include initiatives at the collegiate and K-12 levels to have students delve deeper into applied ethics. MIT's Schwarzman College of Computing has created Social and Ethical Responsibilities in Computing (SERC) case studies[2] for engineering students to explore ethics. Ethics is also at the center of AI4K12's Five Big Ideas in AI framework [33]. Still, few K-12 AI instruction resources discuss ethics with students, and those that do focus on the legal and social implications of deployed or hypothetical systems [98, 99]. The Middle School AI + Ethics Curriculum is an exception to this rule, as it adapts the approach of embedding ethics in technical lessons to develop students' ethical design skills [29, 100]. Doing so enabled middle school students to apply ethical decision making to their AI projects and was highly engaging for the students [100]. One of the Impact.AI

2: MIT SERC case studies, `https://computing.mit.edu/cross-cutting/social-and-ethical-/responsibilities-of-computing/serc-cases-studies` curricula that discussed in Chapter 4, How to Train Your Robot, was built as an extension of the Middle School AI + Ethics Curriculum and similarly prioritizes embedding ethics.

2.3 Designing AI Technologies for Education

2.3.1 A History of AI in Education

According to Chen et al. (2022) [101] the most common applications of AI in education include intelligent tutoring systems for personalized learning, autograders, data mining for performance prediction, computer-supported collaborative learning, social robotics, and affective computing for learner emotion regulation. The most common model for the design of these systems builds on an instructionist philosophy of education where algorithms drive learning, guiding students along some optimized path through the material [102].

Papert (2020) railed against this style of instruction in which "the computer being used to program the child" compromises the learners' creativity and self-determination [48]. Instead, Papert envisioned children programming computers, allowing for immersive learning experiences where students' intrinsic motivation directed their learning.

In considering ways that AI education can support students more holistically, there is promising work in educational social robotics that explores how social robots can boost students' metacognitive skills, such as growth mindset and creativity [103–106]. By employing interaction schema rooted in student-driven pedagogical methods like games, inquiry-based learning, and project-based learning, social robots engage as peer-like learning companions that can support students' learning journeys. This approach starkly contrasts traditional educational approaches where the robot acts as an instructor or coach and directs students to correct answers.

Ali et al. (2021) had students interact with a creativity-eliciting robot that worked alongside them as they completed various drawing and construction tasks [103]. The social robot modeled creativity, prompted users to explore new ideas, and positively reinforced students' creativity. They found that students were twice as creative during creativity-assessment tasks after interacting with the creativity-eliciting robot. A similar example emphasizing the learning benefits of this approach comes from Jung et al. (2014) [107]. Engineering students engaged with a scaffolding agent integrated into an electronics project. The agent asked reflective questions and expressed interest in their creation process. Students working with this agent were more engaged with their task for longer, which led to improved learning outcomes.

2.3.2 Using AI to Teach AI

There is limited research on AI agents scaffolding students' work in constructionist learning environments, aside from the previously mentioned examples. Challenges in this area arise from the unpredictability of open-ended projects, making it hard to predict what support students need. Researchers have gotten around this issue by relying on scripted interactions or utilizing generalized facilitation prompts, like asking reflective questions or presenting contrasting examples, without engaging specifically with students' work [103, 108].

The AI Education scaffolding tools presented in Chapter 6 seek to expand AI's educational potential by creating technologies that holistically support learners and seamlessly integrate into classrooms. The AI education tools discussed in Chapter 6 scaffold students' creativity and technical capability as they create open-ended AI projects. To address some of the ethical concerns with AI in education, myself and the team of researchers developing the tools relied on a human-centered AI development process, collaborating closely with students and educators to inform the design of the technologies.

2.3.3 Challenges Using AI in Education

A significant issue related to educational equity with using AI and robotics in education is the difficulty of bringing new technologies into classrooms. Reich (2020) described a phenomenon called the "EdTech Matthew Effect," where "new technologies disproportionately benefit learners with financial, social, and technical capital to take advantage of new innovations" [102]. Issues around cost, teacher training, and ethical concerns around privacy and students' agency can make it difficult to bring new technologies into classrooms, especially classrooms that are already underserved [102]. Due to these implementation challenges, most new technologies widen the digital divide rather than reduce it, despite technologists' claims about technology and education being great equalizers. In practice, for technology and education to lead to greater equality in opportunities and access, technology creators must active strive to design more inclusive systems.

[102]: Reich (2020), *Failure to disrupt: Why technology alone can't transform education*

2.4 Prior Literature on this Dissertation

Much of the work presented in this dissertation has been previously published in journals and conference proceedings. These published articles include:

A literature review of existing K-12 AI curricula that leveraged the concepts, practices, and perspectives framework to understand What topics AI Literacy researchers typically cover and what assessments they use to evaluate students' learning [109].

Several papers on the RAISE AI Playground and its extensions, including the PoseBlocks extension [69], text classification extension [71], micro:bit robot implementations [110], and Doodlebot extension [111].

Analyses of student performance and teacher feedback on different curricula, including two versions of the How to Train Your Robot (HTTYR) curriculum [110, 112] and Dr. R.O. Bott curriculum [113]. Williams et al. (2022) also described the three design principles for our curricula [112].

Compared to these prior works, this paper offers further analysis of students' performance in activities, the design of the curricula through

[109]: Williams et al. (2023), 'Assessments for K-12 AI literacy: A comprehensive review'

[69]: Jordan et al. (2021), 'PoseBlocks: A Toolkit for Creating (and Dancing) with AI'

[71]: Reddy et al. (2021), 'Text Classification for AI Education.'

[110]: Williams et al. (2021), 'Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum'

[111]: Williams et al. (2024), 'Doodlebot: An Educational Robot for Creativity and AI Literacy'

[113]: Williams et al. (2023), 'Dr. R.O. Bott Will See You Now: Exploring AI for Well-being with Middle School Students'

[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'

the lens of design principles, and discussion on the three curricula in conversation with one another.

The Impact.AI Framework K-12 AI Literacy Framework 3

This chapter introduces the Impact.AI framework, an AI literacy framework that covers the knowledge and perspectives that students should possess to be technosocial change agents. Through the Impact.AI framework, I define a basis to design new K-12 AI curricula and analyze existing curricula. The sections in this chapter are:

- The Impact.AI framework's goals, motivation, and structure
- A comprehensive literature review on K-12 AI curricula that informs the framework
- Details on each component of the Impact.AI framework
- Using the framework to design and review K-12 AI curricula

Pieces of this chapter, particularly the comprehensive literature review and explanation of Impact.AI framework, draw heavily from papers submitted for the author's general examinations. However, translating the Impact.AI framework into learning outcomes and assessments is a significant new contribution of this chapter.

3.1 Overview of the Impact.AI Framework

The mission of Impact.AI is to outline an approach to K-12 AI literacy that equips students to become technosocial change agents capable of leveraging AI to transform society. I correlate students becoming technosocial change agents with them becoming conscientious consumers, ethical engineers, and informed advocates of socially and ethically responsible AI (see Figure 3.1) [114]. In Williams, Ali, Payne, and Breazeal (2019) we discuss how we selected these three goals because we want AI curricula to holistically address the cognitive, pragmatic, and attitude shifts that we envisioned as critical to students becoming empowered users of AI technology [114].

Conscientious consumers. Everyone should be empowered as technically adept AI consumers with the knowledge and confidence to approach new technologies, figure out how they work, and actively consider how they might helpfully integrate new technologies into their lives and redesign technologies to better align with their values.

Ethical engineers. Students should develop a critical lens that allows them to see how AI shapes society and how society can shape AI in return. Whether students shape AI as visionaries, procurers, leaders, or developers, they should have the tools they need to navigate questions like "Should we?" and "In what way?" rather than just "Can we?"

Informed advocates. Referencing the title of this work, which calls for "Democratizing" AI, all people should be empowered to participate in the governance of AI. Today, very few people have the power and access to shape AI. For the vision of democratized AI to be fully realized, all people must understand the political levers they can pull to seek the best

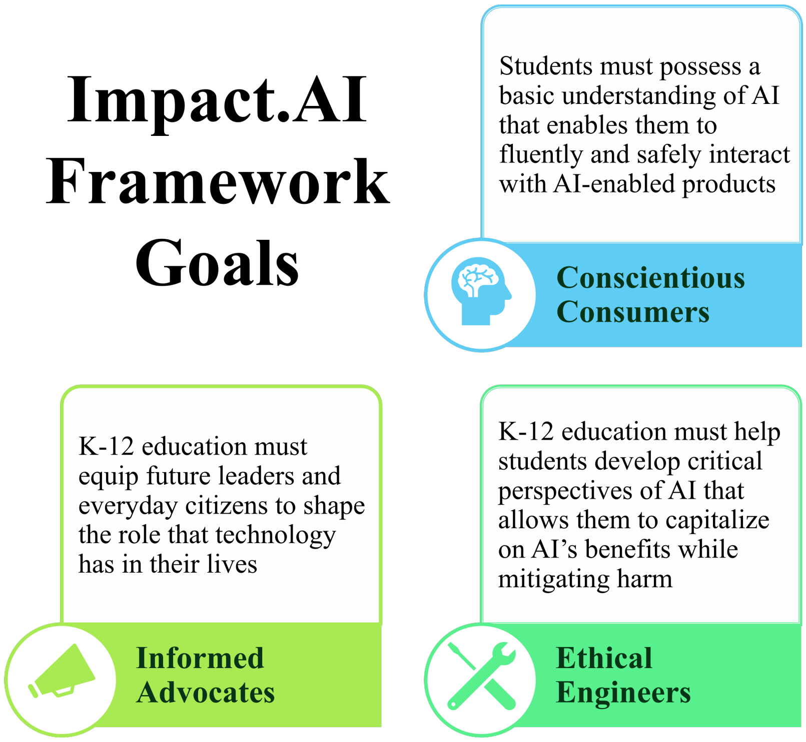

Figure 3.1: Impact.AI framework aims to empower students to become conscientious consumers, ethical engineers, and informed advocates of socially and ethically responsible AI.

for themselves and their communities from the tech industry, including overthrowing AI's present power structure if it is harmful.

3.1.1 Motivation for Impact.AI

The novelty of this framework is that, while defining the AI knowledge and perspectives, the framework also explicitly upholds a justice-centered framing of AI and empowered identities. Only a direct approach to disrupting patterns of harm through responsible AI development and intentional inclusion of marginalized groups will be enough to disrupt the patterns of harm that currently mar the field of AI.

As discussed in the Introduction (Chapter 1, Section 3) and Background (Chapter 2, Section 2), the harms caused by AI systems are significant challenges that the field must address if AI is going to reach its full potential. As an example of how AI harms can be addressed, Suresh and Guttag (2019) describe the machine learning life cycle, from data collection to model development to system deployment, in terms of different kinds of bias and problems that can be introduced at every stage [25]. They emphasize how, regardless of when in the development process an engineer begins and ends working on a problem, both their design decisions and decisions that were made prior and further down the pipeline can have impacts on the end receivers of the deployed system.

Most AI developers are not taught to consider the societal impacts of the systems they develop. Collegiate AI curricula only focus on students developing technical knowledge and skills, neglecting ethics and the societal implications of new systems [92, 93]. Pushing against this harmful trend, Impact.AI centers a community-centered, responsible

AI development process and a pedagogical approach that weaves ethics into technical instruction. In this manner, myself and the educators who participated in this work sought for students to understand that the impact of systems is inextricably linked to how they are designed. Rather than presenting AI algorithms as objective or neutral, from the beginning we pushed students to look for human biases and shortcomings that get baked into systems.

As discussed in the Introduction (Chapter 1, Section 2) and Background (Chapter 2, Section 2) the digital access and participation divide also impacts the field of AI. In their comprehensive literature review on K-12 AI curricula, many researchers have noted the lack of AI curricula specifically focused on diversity and inclusion for demographic groups and entire countries that are left out of and, therefore, underrepresented in the AI workforce [44, 115]. Given the strong relationship between STEM persistence/achievement and STEM identity, bridging the digital participation gap will require us to create learning environments that plant seeds for students to begin developing positive STEM identities [116–118].

While many K-12 frameworks address concerns about broadening participation in AI, only the ML Education Framework and UNESCO Framework for AI Literacy (co-created by Lao, who developed the ML Education Framework) explicitly attend to students' evolving beliefs about AI [37]. In the perspectives category of the Impact.AI framework, the framework specifically attends to students developing beliefs about AI in line with the Culturally Responsive Computing Framework. Myself and the other researchers who contributed to the vision of Impact.AI call for students, especially those who have been excluded and excoded by the tech field, to develop efficacious AI identities where they recognize their power in shaping AI technology and using it for action and expression.

3.1.2 Structure of Impact.AI

The Impact.AI framework is structured in three areas: AI content, practices, and perspectives. Brennan and Resnick's (2012) concepts, practices, and perspectives (CPPs) framework for K-12 computational thinking education inspired this structure [36]. Lao's (2020) ML Education framework, which outlines AI knowledge, skills, and attitudes, also resembles this structure [37]. Impact.AI differs from the computational thinking CPP framework and the ML Education Framework in that Impact.AI focuses on AI-related knowledge for K-12 students.

Given the differences between Impact.AI and these prior frameworks, here I will present exact definitions for how I define the concept, practice, and perspective categories. AI "concepts" are the set of knowledge that AI practitioners engage with as they do their work. AI "practices" are the skills that AI practitioners employ in their work. AI "perspectives" are the beliefs about technology and self that individuals realize as they engage with AI.

Table 3.1: Definition of Responsible AI in the Impact.AI Framework, copied from a white paper by Natalie Lao and Randi Williams (2020)

Responsible	Grapple with issues related to fairness and transparency in the design of AI systems **Encompassed frameworks:** FAccT
Equitable	Evaluate the impact of AI systems on human life, dignity, and rights Design and propose norms and policies to ensure it does
Accountable	Protect end-users by overseeing AI systems and making paths to redress harm **Encompassed frameworks:** Algorithmic Justice
Inclusive	Lower barriers to accessing AI through education and the creation of accessible AI systems **Encompassed frameworks:** Democratizing AI

3.1.3 Defining Responsible AI

Several guidelines around AI ethics have been developed over the past several years as scholarship in the area has evolved with the fast-changing AI field [119]. Thus, as articulated by the title of Fiesler et al. (2020), researchers need to define what "we teach when we teach AI ethics" [92]. Building on top of a whitepaper brief on Ethics in AI Education the author completed with Natalie Lao in 2020 when she and I were part of the Responsible AI for Computational Action (RAICA) curriculum development team, Impact.AI encompasses the principles of fairness, equitability, accountability, and accessibility in AI as described in Table 3.1.

Altogether, these values present a vision for engagement with AI that prioritizes preventing harmful and discriminatory outcomes, setting up systems that protect the agency of those who are impacted by the framework, and designing for accessibility from the beginning of the intervention.

3.2 A Comprehensive Review of Research in K-12 AI Literacy

Before developing the Impact.AI framework, I completed a comprehensive literature review of published research on K-12 AI literacy studies to identify trends and opportunities for future work in the field.

3.2.1 Literature Review Methodology

The search and selection process of the articles included in this review involved using the "snowballing" method outlined by Wohlin (2014) [120]. The snowballing method involves identifying a starting set of papers from relevant research communities and then discovering other relevant papers that cite or are referenced by the starting set. Our starting set included the nine literature review articles mentioned in the Background (Chapter 2), 12 papers from the 2019-2021 proceedings of the Educational Advances in AI (EAAI) and ACM's Special Interest Group on Computer

Science Education (SIGCSE) conferences, and Touretzky (2019) paper [33] on K-12 AI literacy.

One researcher, the author, downloaded and reviewed the abstracts and results of each relevant paper to determine if it met the following criteria:

Criteria 1: The paper is a peer-reviewed article in a conference and journal (i.e., not a thesis, book, or website)
Criteria 2: The paper's educational intervention targets K-12 students
Criteria 3: The paper presents results from a study with K-12 participants.

This work was initially conducted in 2021 and only included papers published until that year, though myself and other contributors to this research have continued collecting papers.

One researcher, the author of this dissertation, completed the data coding and analysis process of this review, the author of this document. The researcher analyzed each included article by recording descriptive information about the articles (the country the authors came from, the year of publication, etc.), then used grounded theory to derive common themes. For the AI content coverage, the researcher categorized the learning objectives described in the paper as an AI concept, practice, or perspective. The search process results are publicly available for future researchers to build upon: https://github.com/randi-c-dubs/k12-ai-ed.

I acknowledge that the coding results are less reliable since only one researcher did the analysis. In the future, a larger team of reviewers will validate this work.

3.2.2 Results of the Literature Review

Before excluding any articles that did not include user studies, I collected 103 papers on distinct AI literacy curricula, proposals, and educational tools for K-12 students published in or before 2021. This collection of more than 100 papers is significant because it is greater than those collected in other literature reviews. This search suggests that more work exists than researchers may know and that work from some countries, authors, and institutions may be systematically overlooked in the review process. Of the papers that met the initial review criteria, 78 conducted a user study with K-12 students. Of these, 63 included the results of the study. These were included in our review.

The next section of text summarizes the articles by the AI content that they portrayed using the concepts, practices, and perspectives (CPPs) framework [36].

[36]: Brennan et al. (2012), 'New frameworks for studying and assessing the development of computational thinking'

AI Concepts

Building of Brennan and Resnick's (2012) [36] examples, I defined AI concepts as the knowledge that AI researchers and practitioners understand and engage with as they work in artificial intelligence, everything from background information about AI to information about how specific

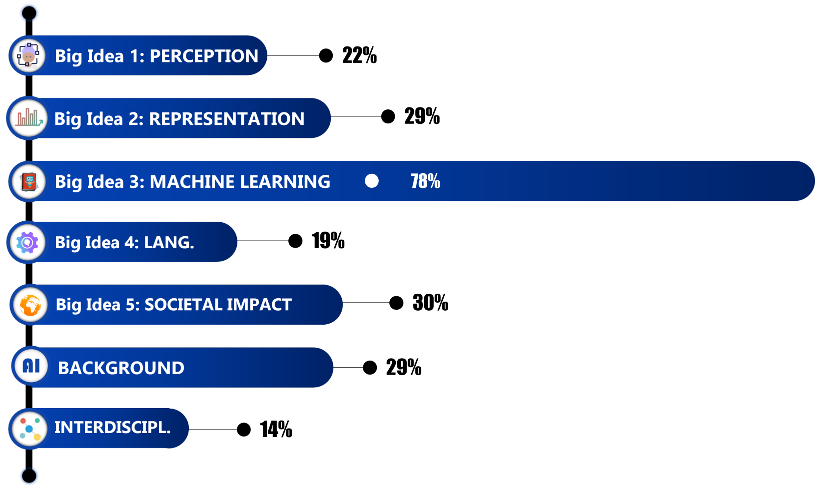

Figure 3.2: The representation of the AI concepts across reviewed articles, n = 63

algorithms operate. Since no taxonomy comprehensively describes all AI knowledge, I categorized knowledge concepts by their alignment with the Five Big Ideas in AI [33], background knowledge, and interdisciplinary knowledge.

[33]: Touretzky et al. (2019), 'Envisioning AI for K-12: What should every child know about AI?'

The Five Big Ideas in AI Are 1. Perception, 2. Representation and Reasoning, 3. Machine Learning, 4. Natural Communication, and 5. Societal Impact. Big Idea #3 Machine Learning was the most covered topic (77.8% of all articles, see Figure 3.2). It was overrepresented compared with the three other ideas that cover similarly essential AI topics. Even within the other three technical areas of the Five Big Ideas, there tended to be a focus on ideas related to Machine learning. Computer Vision was the most explored topic in Big Idea #1 Perception, and curricula about machine learning appeared much more often than other signal processing methods. Data science was often a central theme in curricula about Big Idea #2, and this was to motivate applications in machine learning. Most of the papers about Big Idea #4 Natural Interaction discussed natural language processing and the use of machine learning for intent recognition. Within Big Idea #3 Machine Learning, supervised machine learning applied to image recognition, activity detection, text classification, and similar applications were much more common than curricula about other machine learning topics like reinforcement learning. This focus on supervised machine learning and deep learning likely reflects educators wanting to explain AI topics prevalent in students' lives, namely the supervised ML applications present in many everyday AI technologies.

The next most common topic was Big Idea #5 Societal impact (36.5%). However, they were represented in only about one-third of the published curricula, in contrast to their central position in the Five Big Ideas framework, which suggests they should be present in all curricula. Papers about Societal Impact discussed ethical analysis, ethical design, and the impact of AI systems. Most of these curricula focused on supervised ML and algorithmic bias, but there was some representation of topics like misinformation, ethical theory, and the ethical responsibility of engineers.

Beyond the Five Big Ideas, a handful of papers also taught background information about AI or used an interdisciplinary approach to teach AI. A common theme in these curricula was helping students articulate

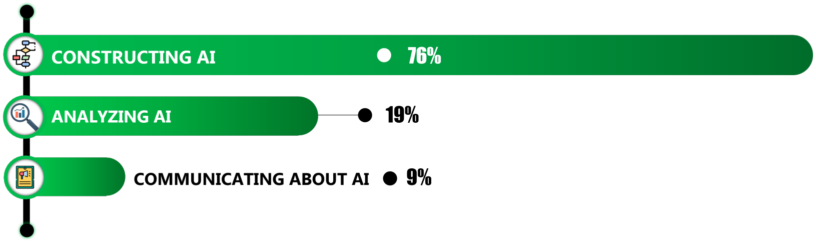

Figure 3.3: The representation of the AI practices across reviewed articles, n = 63

the differences between human and machine intelligence, and the role of biomimicry in algorithm design. An underrepresented topic was the history of AI and how the fundamental ideas in the field emerged. Common interdisciplinary topics included engineering and science, where students could apply AI algorithms. Researchers who leveraged an interdisciplinary approach often leveraged the problem-solving and computational thinking skills inherent to learning about AI algorithms to approach teaching science.

AI Practices

I defined AI practices as the skills and approaches that AI researchers and practitioners employ as they do AI work. From Long and Magerko's AI Literacy framework [34], there were two clear categories of AI practices: constructing AI and analyzing AI systems. To this, we added communicating about AI as our literature review revealed that many curricula included teaching students to collaborate and present their work.

[34]: Long et al. (2020), 'What is AI literacy? Competencies and design considerations'

Constructing AI was by far the most covered practice (73.0% of all articles, see Figure 3.3). The most common AI construction practice was training supervised ML models, which involved using model training platforms that students could use to support these tasks. Graphical and tangible user interfaces were commonly used to support younger learners [131–133, 147, 162, 167] and those who are new to computer science [67]. Amongst these, the validation and testing steps were less commonly taught than the training portion of the model creation process, indicating a decreased priority in teaching students how to evaluate and assess their models.

Fifteen (15), 23.8% interventions included programming or computational thinking as a practice. Block-based programming languages were often used to make AI projects easier for students, even if they did not have a background in programming. Many of these systems build on popular, well-established programming platforms like Scratch [68, 69, 134, 138, 173], AppInventor [166], and LEGO Mindstorms [124], and Snap [55, 129]. Researchers used text-based languages to bridge the gap between the classroom and expert practice. Mobasher et al. had students explore data science using field-appropriate tools like Matlab, SPSS, and Tableau [136]. Similarly, Mariescu-Istodor had students bring their AI application to life using HTML and JavaScript [135].

After practices related to constructing AI artifacts, the next most popular set of practices contained ways of analyzing AI artifacts and was present in eighteen (18), 28.6% interventions. Two kinds of analysis, technical and ethical, were present in existing education interventions. Technical analysis skills included having students dissect AI systems' functionality by identifying major algorithms' major components [58–60] or predicting what kind of training data was used [29]. Ethical analysis skills included having students consider ethics throughout all stages of an AI system - design, implementation, and deployment. A set of researchers from the MIT Media Lab taught middle school students about stakeholder analysis [29, 100, 110], while other papers had students evaluate the presence of bias in models [29, 160, 168] and to consider the ethical implications of systems [60, 142, 152], including systems that students designed themselves [69, 110].

The least commonly occurring set of practices included those about communicating about AI. This set of important practices included helping students share their work using science communication standards [137], collaboration skills [125, 148], and presenting final projects [125, 136]. Forsyth et al. (2021) incorporated storytelling and media creation lessons to teach students to communicate about current issues in AI [152].

AI Perspectives

Finally, we defined AI perspectives as the beliefs and evolving understandings students develop as they learn about AI. In existing interventions, we separate these perspectives into two categories: students' point of view as responsible AI stewards and self-awareness of their role in an AI-driven society.

Since perspectives and attitudes are not part of most AI literacy frameworks, we saw interventions aimed toward these perspectives less often. Some curricula involved students first identifying AI artifacts [60, 146, 172], then becoming more aware of the limits of AI and that it can make mistakes [129, 146, 165], and then seeing how much human subjectivity could influence those AI biases and mistakes [153, 174]. Other important ethical perspectives for responsible AI included having students see that AI systems can have both positive and negative effects [59, 60] and that often, the impacts of these systems are felt differently by different stakeholders [29, 60, 165].

As for students' AI identities, major themes included helping students develop self-efficacy in AI [166, 172, 175]. A few curricula explicitly focused on students from historically excluded communities and developing a sense of belonging for their students by introducing them to role models [125, 148, 165], helping them build a sense of community [165], and exposing them to AI careers [60].

3.2.3 Review of K-12 AI Literacy Work Published Between 2022-2023

Although the literature review concluded in 2021, research in K-12 AI literacy has continued rapidly and resources such as Google Scholar and key conferences like CHI and EAAI have made it possible to continue collecting papers. Continuing the growth pattern seen in previous years, I have collected 85 additional papers published between 2022-2023 that may meet the inclusion criteria of the literature review. Of those, 38 were published in 2022, and 47 were published in 2023. This section includes a cursory overview of how issues of inclusion, ethics, and teacher preparation have played a more prominent role in recently published articles.

A large chunk of recent work, 15 out of the 85 collected papers, has been done with and for K-12 teachers. One branch of this work surveyed teachers. They found that, although teachers were interested in bringing AI to their classrooms, they had pressing needs to obtain pedagogical content knowledge and have space in their school day to engage with AI. Because of their limited knowledge, researchers observed that some teachers had misconceptions or fears about AI that limited their teaching [176–178]. Other researchers focused on structural issues like teachers not being supported by their schools to teach AI [178]. Partnerships between AI education researchers and teachers developing AI curricula are one way to facilitate AI instruction in their classrooms [179, 180]. Overall, the findings of surveys with teachers pointed to the importance of schools adopting standards to promote AI instruction and curriculum developers creating appropriate curricula [176].

Many researchers have launched co-design collaborations [181, 182] and professional development programs [179–181, 183–186] with instructors. Amplo et al. (2023) highlighted the importance of profound collaborations with teachers where they played an active role in shaping the curriculum, even as they gained more knowledge about the field [181]. Many partnerships empowered teachers as creators, not just implementers, of AI curricula [179, 184, 185]. Researchers provided teachers with extended professional development experiences, including teaching practica and community-building opportunities [178, 185, 186]. Cu et al. (2023) proposed a design for a new application that assists teachers in making lesson plans [183].

A notable change in the landscape of AI since 2021 is the emergence of ChatGPT as a consumer application and educational tool. Although none of the 85 articles explicitly cover generative AI, many initiatives to teach K-12 students and teachers about AI have emerged. Notable examples include the lessons on ChatGPT created by Code.org for their 2023 Hour of Code initiative [187] and the ChatGPT in School lesson created for the 2023 Day of AI outreach event [188]. These resources engage experts in explaining the inner workings of ChatGPT and other large language models and guide students to think through appropriate use cases, limitations, and potential harms of using these tools. On top of these resources for students, resources and tools for teachers to learn about and leverage ChatGPT and other AI tools in their classrooms have also hit center stage. The speed with which these resources were developed underscores how quickly the field of K-12 AI literacy moves. In future curricula, the popularity of generative AI will likely increase the overrepresentation of curricula about machine learning but also boost the representation of curricula about chatbots and natural interaction with AI. Researchers should aim to develop curricula that balance topics to help students maintain a well-rounded perspective on the subject.

Research efforts and curricula addressing issues of inclusion and ethics are rising in recently published curricula. In papers published from 2022-2023, 10 articles explicitly addressed ethics and critical digital literacy, two addressed broadening participation in AI, and another 10 addressed some aspects of both.

In their surveys about teachers capabilities to teach AI, Polak et al. (2022) [189] and Gibellini et al. (2023) [190] note that it can be challenging to design for inclusion because of the broad range of prior experiences with AI that students bring to the classroom. Particularly for research efforts taking place outside of the Western world, including Martins et al. (2023) [191] in Brazil and Sanusi et al. (2023) [192] and [193] in Nigeria, issues of accessibility and contextualizing curricula in the cultural and experiential knowledge of learners has been critical to the success of their lessons.

A handful of researchers explicitly leveraged culturally responsive-sustaining pedagogy [87, 194–197] and family and community engagement [198, 199] to ensure AI topics were relevant. For example, Lee et al. (2022) engaged youth in learning about AI by having them research and create content about the intersections of AI and social justice issues in their communities [195]. Researchers also engaged in AI instruction to grow students' sense of AI self-efficacy [87, 197, 199, 200]. The increased focus on issues of inclusion for learners beyond dominant groups reflects and strengthens the expanding impact of AI on different members of global society.

In their survey on teachers' perspectives on teaching AI, Gibellini et al. (2023) [190] also observed that teachers had mixed reactions to teaching AI. They believed that the topic was important but needed help to teach it compellingly. Researchers addressing ethics in their curricula covered a range of topics. The most common ethics topics were around the real-world societal impact of AI [182, 194, 201–204]. Other researchers highlighted specific harms, such as algorithmic bias, or values, such as fairness and accountability in AI [87, 196, 200, 205]. Researchers also discussed the limitations of AI [194, 206], the subjectivity of their creation [202, 206], and bigger picture ethical dilemmas [193, 207]. From AR [201] simulations to museum exhibitions [207] to discussions, researchers leveraged various strategies to have actionable and perspective-changing conversations about AI.

3.2.4 Literature Review Discussion

Completing the literature review led to four broad observations about the current state of K-12 AI literacy research and future work to push the field forward.

The Challenges of Teaching Students About a Fast-Paced Field

Seeing the exponential growth in published articles offers a sense of perspective about the fast pace of this work. As AI rapidly evolves, AI instruction will need to keep pace. This is much different from other fields, such as literature arts and mathematics, where innovation does not drive what children need to know. K-12 AI literacy researchers must latch on to the fundamentals of AI, including the history of AI and design processes that can last over time even as they strive to remain nimble and responsive to new technologies.

Developing Educators' Technical, Pedagogical, and AI Content Knowledge

It will be increasingly important to bring K-12 AI curricula to classrooms and get educators involved as much as possible. Most existing resources for K-12 AI education are generally short, informal learning opportunities focusing on a particular concept or big idea in AI. However, as De la Higuera (2019) argues, schools should make time to teach AI in their classrooms, given the relevance of AI in children's lives [208]. Broadening the reach of AI instructions means training K-12 educators to teach AI, though there are few published pieces on preparing in-service teachers to bring AI education to their classrooms. For AI curricula to move from research studies to the classroom, pedagogical approaches must be practical and easy to learn, and programs must be developed to equip teachers with knowledge and resources to support their students.

Teaching What Matters in K-12 AI

The content of AI courses must match what is important for students to learn. Machine learning is overtaught at the expense of other topics, particularly societal and ethical lessons that tend to be overlooked in collegiate AI courses [92]. Suppose societal impact and ethics are to be the center of learning about AI, as it is presented in the Five Big Ideas in AI [33]. In that case, researchers must prioritize teaching AI ethics in their courses. Also, as trends in AI move, researchers should strive to maintain a balance between teaching the newest innovations and ensuring students get a comprehensive of the entire field so that they can adapt when another wave of innovations occurs.

Broadening Participation in K-12 AI Literacy Research

We saw that most published work in this area comes from universities in the United States. Just as the creators of AI need to be diverse, the creators of AI curricula should be as well. Otherwise, potential barriers to broader adoption of curricula will exacerbate the digital divide. The technological feasibility and cultural relevance of curricula greatly depend on curriculum creators' deep knowledge of the students they are trying to serve.

Table 3.2: Reviewed articles by year of publication

Year	N	Studies
2008	1	Bigham et al.
2011	1	Rosen, Silverman, and Essinger
2014	1	Benotti, Martínez, and Schapachnik
2016	3	Kandlhofer et al., Burgsteiner, Kandlhofer, and Steinbauer, Vachovsky et al.
2017	1	Srikant and Aggarwal
2018	4	Hitron et al., Sakulkueakulsuk et al., Kahn et al., Ureta and Rivera
2019	10	Druga et al., Hitron et al., Williams, Park, and Breazeal, Williams et al. Tang, Utsumi, and Lao, Zimmermann-Niefield et al., Estevez, Garate, and Graña, Mariescu-Istodor and Jormanainen, Mobasher et al., Zhang et al.
2020	13	DiPaola, Payne, and Breazeal, Alturayeif, Alturaief, and Alhathloul, Bilstrup, Kaspersen, and Petersen, Lin et al. Norouzi, Chaturvedi, and Rutledge, Sabuncuoglu, Schaper, Malinverni, and Valero, Shamir and Levin Skinner, Brown, and Walsh, Van Brummelen et al., Vartiainen, Tedre, and Valtonen Vartiainen et al., Wan et al.
2021	29	Ali, DiPaola, and Breazeal, Ali et al., Ali et al., Choi and Park, Druga and Ko Lee et al., Jordan et al., Forsyth et al., Henry, Hernalesteen, and Collard, Kaspersen, Bilstrup, and Petersen, Lin et al., Kaspersen et al., Kim et al., Long et al., Lyu, Ali, and Breazeal Melsión et al., Olari, Cvejoski, and Eide, Park et al., Rodríguez-García et al. Tseng et al., Shamir and Levin, Aki Tamashiro et al. Van Brummelen, Tabunshchyk, and Heng, Vartiainen et al., Voulgari et al. Williams, Yoder et al., Zhang et al., Zhu and Van Brummelen
Total	63	

Factor	Category	N	Percent
Age group	Primary	11	17%
	Middle	11	17%
	Secondary	22	34%
	Pre-K and Primary	2	3%
	Primary and Middle	6	9%
	Middle and Secondary	4	6%
	Primary to Secondary	6	9%
	Primary to Adult	2	3%
Setting	Informal	40	63%
	Formal	13	20%
	Laboratory	11	17%
Country of first author	USA	34	53%
	Denmark	4	6%
	Finland	4	6%
	Israel	4	6%
	Spain	3	5%
	Austria	2	3%
	Korea	2	3%
	Other	11	17%

Table 3.3: Attributes of reviewed articles

3.3 The Impact.AI Framework

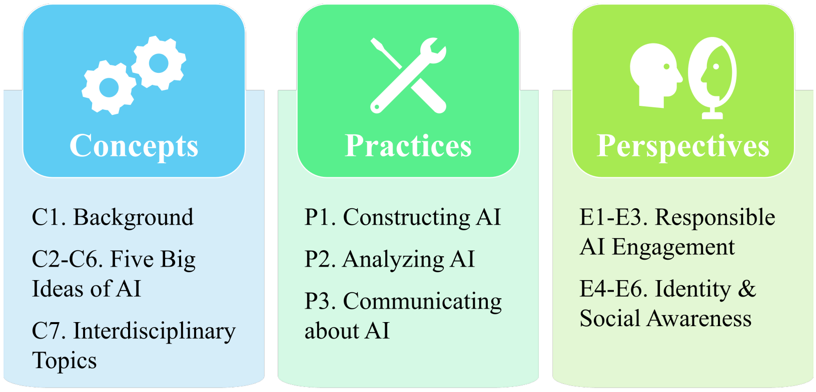

Figure 3.4: The Impact.AI Framework featuring its three categories: AI concepts, practices, and perspectives

We defined AI concepts based on the Five Big Ideas in AI [33] and the 17 AI competencies [34] because the two frameworks provide a comprehensive view of AI knowledge. First, students need background knowledge about AI, including what AI is and what it can do [34]. Next, they must understand the Five Big Ideas in AI: perception, representation and reasoning, learning, natural communication, and societal impact [33]. Finally, students should be taught interdisciplinary topics (e.g., biological science, data science, math, math) alongside AI to connect ideas to other subjects.

For Big Ideas #1-4, we derive our expectations for student knowledge from the Five Big Ideas in AI [33] and Long and Magerko's (2020) [34] AI literacy competencies. The goal is for students to articulate and understand key ideas. For example, they should be able to clearly define each Big Idea, list several applications in each area, and know about and be able to simulate how specific algorithms work. For interdisciplinary knowledge of AI, we take our cue from Long and Magerko's AI competencies definition, which articulates that students should be able to use their AI knowledge to understand other subjects and vice versa [34]. For all of these categories, we imagine a particular curriculum, based on their target age range and the length of the intervention, to select one or two main AI or interdisciplinary topics as a focus.

Background information about AI and Big Idea #5 Societal Impact are slightly different. Drawing from Long and Magerko (2020) [34], every student should have a working definition of AI that articulates the different kinds of intelligence referred to by the term "AI" and can expand as the field evolves. Plus, students should have some historical context for AI, who the main players were, and what their aims were. In Societal Impact, students should know several common benefits and harms of AI, like those presented in the AAAI 20-Year community roadmap for AI research [209], be able to describe different perspectives about AI, and be able to articulate how humans play a role in AI development, for example by labeling diagrams of the AI development pipeline such as was done by Suresh and Guttag (2019) [25].

Next, practices are defined by Brennan and Resnick (2012) as the skills and methods that AI practitioners employ in their work [36]. AI construction skills are commonly taught in prior work – many curricula provide training, programming, and prototyping tools. Another component of skills related to Constructing AI that should be included are that students should be able to determine to what extent a problem can be addressed with AI and they should be able to use other methods throughout the community-centered design process in their development of AI systems.

Analysis and communication appeared less often although they are important skills to emphasize if we are to build AI systems that protect the most vulnerable members of society [29, 31, 100]. To prepare students to critically examine the societal consequences of technology, analyzing and communicating about AI must be more central in AI literacy.

In analyzing AI, we intend for students to be able to analyze AI as a sociotechnical systems. Students should be equipped to employ both technical and ethical standards and methods to evaluate the design, functionality, and impact of AI systems. Some tools that come to mind include methods like using testing datasets to evaluate the accuracy of supervised machine learning systems, including the false positives, false negatives, and intersectional analysis on different groups within the test data set. Ethical analysis skills include conducting stakeholder analyses to understand the values and potential impacts of systems on different groups who will be impacted by a system's design. Or using impact analyses to rate the risk level and scope of AI systems, and design appropriate mitigation for the system.

Critical communication tools include students learning how to collaborate with interdisciplinary groups to develop AI systems. This will require students to learn how to communicate about the design, technical construction, and impact of systems on audiences at different levels. Two powerful tools they could learn regarding communicating with different audiences come from Buolamwini et al. (2022) who describes the different powers of using algorithmic vs. evocative audits [30]. Algorithmic audits are systematic analyses of technological systems that usually involve using technical analysis methods to describe the accuracy of a system. Evocative audits are used to humanize the impact of a deployed system, for example, by using a "counterdemo" that clearly shows a deployed system failing in some critical way in a situation that it claimed to function well. Knowing how to leverage both of these tools is a powerful way to communicate about algorithms across audiences with different educational backgrounds and tools for redressing potential harms.

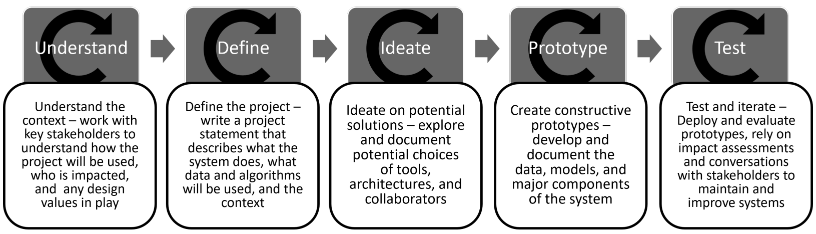

Figure 3.5: The community-centered AI design process as described in the SIMPL ML Model project

Finally, perspectives are the beliefs about technology and self that individuals realize as they engage with AI. Our framing of AI perspectives builds on the "Three Practices of 21st Century Citizens" model from Torney-Purta et al. 2004 [210] and the Culturally Responsive Computing framework by Scott et al. (2015) [26]. From the "21st Century Citizens" model, we derived three areas of beliefs around technology's role in shaping society and society's role in shaping technology.

Torney-Purta et al. (2004) developed a citizenship framework for K-12 students that outlines competencies for students that contribute to their civic engagement [210]. Like Brennan and Resnick (2012), this framework also addresses concepts, practices, and perspectives, framed as knowledge, skills, and dispositions, where dispositions are motivations for behavior, values, and attitudes [36]. Torney-Purta et al. (2004) describe dispositions for civic engagement in terms of civic literacy and digital citizenship for informed, engaged, and active citizenship [210]. From these ideas, we derived AI perspectives around responsible AI engagement, where students should develop attitudes that signal digital literacy, critical digital literacy, and digital citizenship.

[210]: Torney-Purta et al. (2004), 'Developing Citizenship Competencies from Kindergarten through Grade 12: A Background Paper for Policymakers and Educators, Education Commission of the States, Denver, Colorado, USA'

The first area, digital literacy, includes awareness of AI's impact on students' personal lives, culture, and future careers. In this perspective, students are more generally aware of the role of AI in their individual lives and community. With this awareness, students should become more curious about the behind-the-scenes roles of AI in their lives, more engaged with using AI as a tool, and equipped to navigate AI's presence with agency. The second area, critical digital literacy, involves students' awareness of the limitations of AI and different perspectives about how AI might be used. By recognizing AI's limitations, students should be able to use their knowledge of the societal impact of AI and their skills in analyzing AI to critique new AI systems that become part of their lives. They should have a balanced view of these systems' potential benefits and harms. Students should be able to push against the "techno solutionist" and "techno-optimist" perspectives that are typically embedded in computing education. By recognizing that many different perspectives about AI exist, they should welcome hearing new perspectives and be prepared to use their AI communication skills to share their own. Finally, as engaged digital citizens, a notion that we extend to all students regardless of their nationality or residential status, students should be aware of and engaged in the democratic regulation of AI. When students feel strongly that their involvement in AI governance matters, whether they see themselves as AI developers or not, they may take action to improve the field.

From the five tenets of the "Culturally Responsive Computing" framework, we derived three areas encompassing students' understandings of themselves and their potency as tech activists. The first area of critical importance for identity and social awareness is that students recognize that they, themselves, and all people, are capable of digital innovation. Students should form a STEM identity and have foundational skills to rebut harmful stereotypes about who can and cannot be successful in the tech field. The second area of importance draws from the second and third tenets of Culturally Responsive Computing - that students know they can use technology for expression and action. This point strongly resonates with constructionism and computational action, which was discussed in Chapter 2. The third and final area of identity and social awareness draws from the fourth and fifth tenets of Culturally Responsive Computing, that students should measure success by how much they use their knowledge to further their communities and as a means toward a social justice end. This area involves students developing a community and sense of belonging in the field of AI and collaborating with robust communities full of different skill sets. These areas of perspectives are particularly critical to include and encode all students.

3.3.1 Designing Curricula with the Impact.AI Framework

Learning Outcomes

One purpose of the Impact.AI framework is to support curriculum designers and educators in implementing effective AI lessons. Impact.AI can inform the learning outcomes and correlating assessments to organize and evaluate new AI curricula. An excerpt of the learning outcomes from the Appendix is shown in Table 3.7.

Assessments

> “measure what you value instead of valuing only what you can measure.”
>
> Andy Hargreaves,
> 2011 International Confederation of Principals

In Williams and Breazeal (2023) we published a literature review of K-12 AI curricula explicitly focusing on the assessments used in each intervention [109]. Building on this prior work, this section presents a specific approach to comprehensively assessing students’ progress toward Impact.AI’s academic goals.

[109]: Williams et al. (2023), ‘Assessments for K-12 AI literacy: A comprehensive review’

Summarizing the results of that literature review, we found many assessments of students’ knowledge and AI construction skills, particularly in machine learning. However, there are no standardized evaluations, and few have been validated, meaning that it is difficult to compare the results of different curricula or students’ learning across contexts or timespans. It is of the utmost importance that validated and robust assessments for AI comprehension are developed so that we can move toward best practices in the field. Below, we will discuss specific recommendations for assessing AI concepts, practices, and perspectives.

Written assessments were the most common approaches for evaluating students’ understanding of AI concepts. Two kinds of written assessments can be beneficial: concept inventories created by researchers about specific topics and self-assessment questions. Rodriguez-Garcia et al. (2021) [163] and Lee et al. (2021) [60] both developed two AI concept inventories that assess background AI knowledge (define AI) and topic-specific (what is true of supervised machine learning) knowledge that students should understand. Both of these concept inventories mainly focus on machine learning knowledge, so there is an opportunity for future assessments to cover a broader range of knowledge. Self-assessment questions were also commonly used in interventions. Although they are less objective, self-assessments can be helpful for students to have a metric on which to measure their knowledge growth both formatively and summatively as they advance through a curriculum.

Table 3.4: The seven categories of AI concepts in Impact.AI

AI Concept	Examples
C1. Background	What is AI, history of AI, human vs. machine intelligence
The Five Big Ideas In AI	
C2. Big Idea #1: Perception	Sensors, computer vision, signal processing
C3. Big Idea #2: Representation and Reasoning	Automata, data structures, search, path planning
C4. Big Idea #3: Machine Learning	Unsupervised machine learning, reinforcement learning, data, classification, generative models, transfer learning
C5. Big Idea #4: Natural Communication	Speech synthesis, chatbots, autonomous vehicles, human-computer interaction
C6. Big Idea #5: Societal Impact	Ethics, design values, environmental impact
C7. Interdisciplinary Topics	Sustainability, aeronautics, bioinformatics, robotics, ecosystems

Table 3.5: The three categories of AI practices in Impact.AI

AI Practice	Examples
P1. Constructing AI	Problem scoping, design thinking, creating ML and non-ML models, programming
P2. Analyzing AI	Discerning the inputs and output of ML systems, identifying stakeholders, bias assessments
P3. Communicating About AI	Storytelling, technical writing, creating persuasive arguments, advocacy, collaboration

Table 3.6: The six categories of AI perspectives in Impact.AI

AI Perspective	Examples
Responsible AI Engagement	
E1. Digital Literacy	Awareness of AI in personal life, awareness of AI in future careers, awareness of AI's impact on culture
E2. Critical Digital Literacy	Recognizing different perspectives, recognizing that AI systems are programmable
E3. Digital Citizenship	Valuing different voices and contributions to AI, participating in democratic regulation of AI
Identity and Social Awareness	
E4. Self-Efficacy	Belief in one's capability
E5. Activism and Expression	Using AI to take meaningful action, aware that AI can be used to express oneself
E6. Community	Recognize self as part of a larger community, identify with AI role models

Table 3.7: Sample learning outcomes for the Impact.AI concepts. See the rest of the Impact.AI learning outcomes at `https://tinyurl.com/impactai-objectives`.

Concepts	
In demonstrating their understanding of fundamental concepts in AI, students will be able to:	
C1. Background C2. Big Idea #1 C3. Big Idea #2 C4. Big Idea #3 C5. Big Idea #4	Provide an accurate and precise definition of AI and key terms related to the Big Ideas Articulate differences between natural and machine intelligence by listing them and drawing comparison diagrams Describe how different AI algorithms work by creating or labeling high-level process diagrams of the algorithms' inputs, outputs, and main components Describe several examples of technologies that use AI and articulate the capabilities and limitations of the AI systems
C6. Big Idea #5	Describe several ways that AI systems generally benefit and harm society by listing them Describe different perspectives on the present and future impact of AI by creating descriptive diagrams Articulate how humans play a role in every part of the AI development process by creating high-level process diagrams
C7. Interdisciplinary Topics	Articulate the connections between interdisciplinary subjects and AI by creating comparison diagrams Describe several examples of systems at the intersection of AI and the interdisciplinary subject

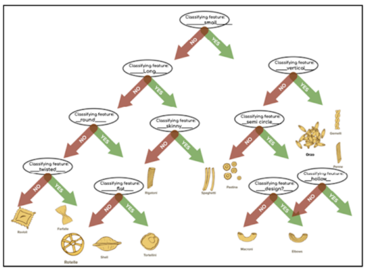

Figure 3.6: Screenshot of a diagram from the AI-Concept Inventory developed by Lee et al. (2021) [60]

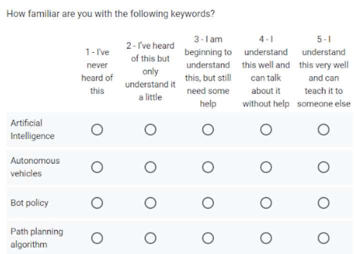

How familiar are you with the following keywords?

	1 - I've never heard of this	2 - I've heard of this but only understand it a little	3 - I am beginning to understand this, but still need some help	4 - I understand this well and can talk about it without help	5 - I understand this very well and can teach it to someone else
Artificial Intelligence	○	○	○	○	○
Autonomous vehicles	○	○	○	○	○
Bot policy	○	○	○	○	○
Path planning algorithm	○	○	○	○	○

Figure 3.7: Screenshot of a knowledge self-assessment question from Williams, Ali et al. (2024) [113]

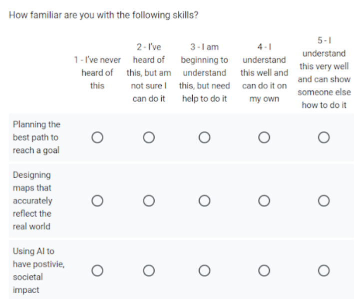

How familiar are you with the following skills?

	1 - I've never heard of this	2 - I've heard of this, but am not sure I can do it	3 - I am beginning to understand this, but need help to do it	4 - I understand this well and can do it on my own	5 - I understand this very well and can show someone else how to do it
Planning the best path to reach a goal	○	○	○	○	○
Designing maps that accurately reflect the real world	○	○	○	○	○
Using AI to have postivie, societal impact	○	○	○	○	○

Figure 3.8: Screenshot of a skills self-assessment question from Williams, Ali et al. (2024) [113]

[60]: Lee et al. (2021), 'Developing Middle School Students' AI Literacy'

[211]: Carolus et al. (2023), 'MAILS–Meta AI Literacy Scale: Development and Testing of an AI Literacy Questionnaire Based on Well-Founded Competency Models and Psychological Change- and Meta-Competencies'

Another helpful tool for assessment of AI concepts could be concept discussions and artifact reviews. Brennan and Resnick (2012) [36] used this idea in their assessment of computational thinking education, and it has also found its place in K-12 AI literacy [66, 164]. An artifact interview consists of students thinking aloud as they work on projects or reflecting on the projects they have created, leveraging their conceptual knowledge to explain different concepts of their work and how it works. The benefit of this approach is that it allows students to ground their knowledge in the narrative of their project creation, allowing for a more subjective evaluation of their knowledge.

Similar tools can come into play for evaluating students' knowledge of AI practices. Skill inventories and artifact interviews can help educators measure how much students have mastered essential skills. Sakulkueakulsuk et al. (2018) had students use technical performance metrics to have teams compete against one another to create projects with the best testing performance since the intervention was about supervised machine learning [128]. In Williams et al. (2022), we used a combination of technical design, ethical design, and implementation quality to evaluate students' projects [112]. We graded students' mastery of technical and ethical analysis skills based on their submitted project code and documentation.

A final option for assessing students' skills is using activity-based assessments where students are tested on specific skills as they work through learning activities. For example, Zhang et al. (2019) gave students logic problems they had yet to see in a formative unit assessment [137]. DiPaola et al. (2020) described the results of a design problem where students used ethical thinking practices like stakeholder analysis on case studies [100].

Finally, there are three methods to evaluate students' perspectives: surveys, discussions, and activity-based assessments. Existing surveys for AI knowledge typically cover students' perceptions of AI and attitudes toward AI. Lee et al. (2021) validated an Attitude Toward A Questionnaire and AI Anxiety Scale that measure students' interest in AI and their perceptions of the dangers of AI [60]. Similarly, Carolus and Koch et al. (2023) [211] produced a MAILS Meta AI Literacy Scale that measures AI literacy and self-efficacy. These tools can do an excellent job of measuring students' digital literacy, self-efficacy, sense of belonging to the community, and attitudes toward AI. However, they fail to evaluate critical digital literacy and aspects of students' technosocial change agent identity.

Further tools should be developed to create a more comprehensive view of students' attitudes. A tool such as the "Draw a Scientist" activity may be used to understand students' changing understanding of stereotypes about who can be successful in AI [212]. This activity, originally developed to assess students' experience of traditional gender stereotypes in science, has since been used to unearth other stereotypes students may hold about race, gender, class, and more in other STEM fields.

Discussions and debates can help students clarify their beliefs about AI and make stronger arguments to support their beliefs. Discussions allow students to give and receive feedback won different perspectives about AI, particularly related to digital and critical digital literacy. Von Wangenheim (2020) had students debate different ethical issues in AI [173]. These debates revealed students' beliefs about AI, particularly their beliefs about appropriate roles for AI, given the limitations of technology. Activity-based assessments of perspectives allow students to develop their understandings and put them into action. Ali et al. (2021) had students put their opinions about AI into action using an activity where students advocated for different policies for regulating AI [59]. In this way, they allowed students to articulate their opinions about AI while also considering actions they might take to become AI activists. Activity-based assessments allow students to express and take action on their ideas.

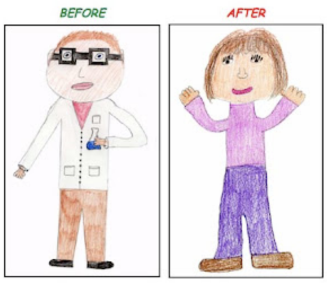

Figure 3.9: Image of the Draw a Scientist activity from a STEM outreach activity at Fermilab, Illinois, USA. 2010.

In Chapters 4 and 5, I will describe two middle school AI curricula I developed using the Impact.AI standard as a basis. The first curriculum, How to Train Your Robot, came before the development of the Impact.AI framework, and much of its construction informed the creation of the framework and the assessment recommendations presented in this section. AI for Wellbeing came after the development of the Impact.AI framework, and its learning outcomes and assessments originated from the recommendations in this section.

3.3.2 Analyzing curricula with the Impact.AI Framework

I converted the Impact.AI Framework into a rubric for evaluating AI curricula. Like the ML Education Framework Rubric [37], curriculum creators and educators can use this tool as they design and select resources for their students. This section presents the Impact.AI rubric and summarizes the evaluation of three existing AI curricula.

[37]: Lao (2020), 'Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens'

The rubric contains sixteen items, seven areas of concepts, three areas of practice, and six areas of perspectives, just like the Impact.AI Framework. For each item, a curriculum can receive a score of 0 to 4 with a "0" or no coverage, meaning that the curriculum does not meet the criteria for a "1" or minimal coverage. A "4" means that there is advanced coverage of that particular item.

A single course or curriculum is not expected to provide advanced coverage of all or even most items on the rubric. Considerations such as the prior experience of the intended audience, time constraints, and resource constraints make it difficult for educators and curriculum designers to cover all content at the advanced level. That said, the rubric is beneficial as it indicates to educators whether a particular resource covers the content they believe will be most helpful to their students.

In the general exam submitted in 2021, I used the Impact.AI rubric to evaluate three AI curricula: the DAILy curriculum, ISTE's Hands-On AI Projects guide, and Elements of AI. The DAILy curriculum is for middle and secondary teachers and was created in collaboration with the MIT STEP Lab, MIT Personal Robots Group, and Boston University. The curriculum contains 30+ hours of content to introduce students to AI concepts, ethical issues in AI, and the presence of AI in future careers [60]. ISTE's Hands-On AI Projects guide for secondary school teachers contains about 30 hours of student-driven AI projects to explore the fundamentals of AI technologies [213]. Finally, Elements of AI is a massive open online course (MOOC) created by MinnaLearn and the University of Helsinki. It seeks to reach a broad, public audience and teach them the fundamentals of AI [214].

[60]: Lee et al. (2021), 'Developing Middle School Students' AI Literacy'

[213]: ISTE (2022), 'Hands-On AI Projects for the Classroom: A Guide for Secondary Teachers'

[214]: University of Helsinki and MinnaLearn (2021), 'Elements of AI: Introduction to AI'

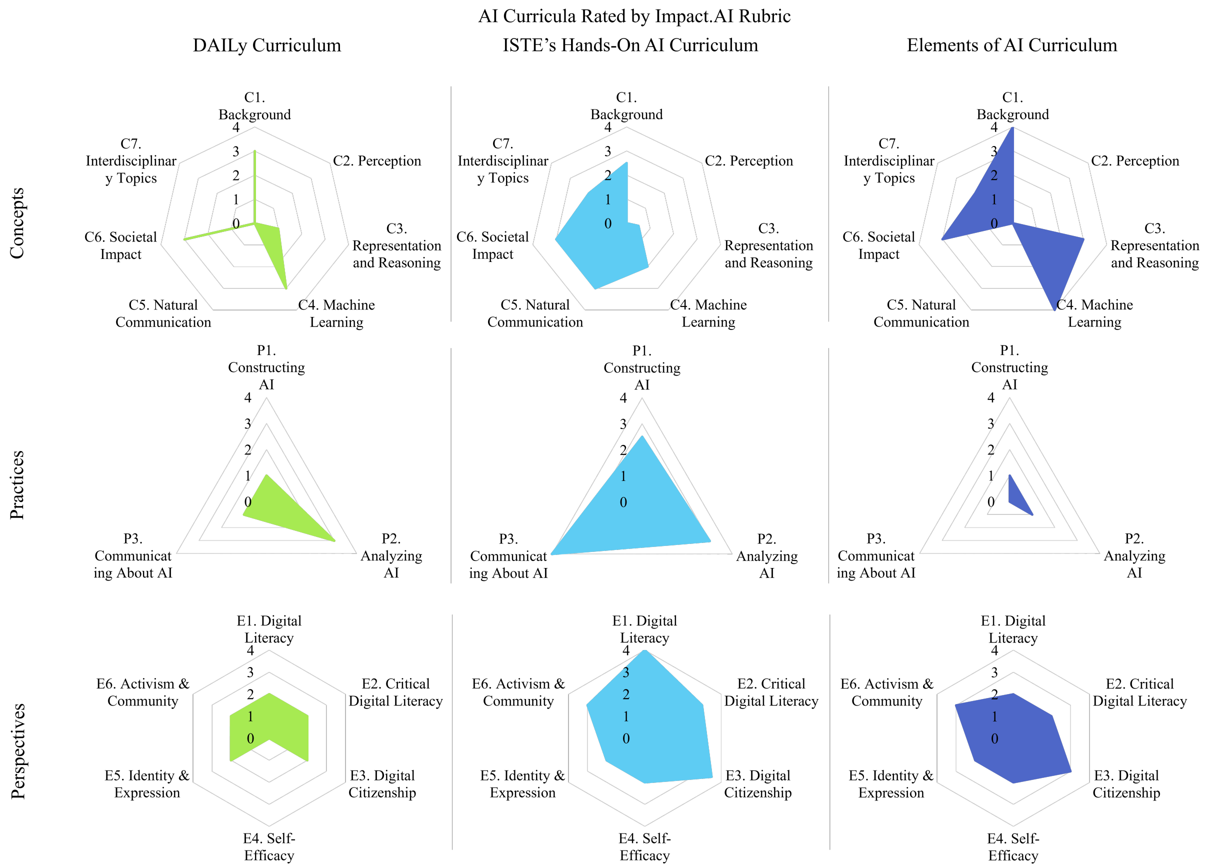

Figure 3.10: Side-by-side comparisons of the Impact.AI rubric scores of three AI curricula: the DAILy curriculum, ISTE's Hands-On AI Projects guide, and Elements of AI.

The results of that analysis highlighted how differences in motivation shaped the content included in each curriculum. Although many curricula strove to educate people from diverse audiences, all fell short of designing culturally relevant, responsive, and revitalizing curricula to empower students. AI curricula must strive toward this goal to address the systemic barriers preventing some people from entering and thriving in AI.

Another observation was that every curriculum included a different set of AI content, depending on the length of the intervention and the year the curriculum was published. These differences are related to the rapid change of AI, implying that the most important lessons to impart should prepare students for future learning and engagement with AI. Rather than focusing on any particular topic, every AI curriculum should help learners develop skills to continue seeking and building their knowledge.

Finally, AI curricula should be more intentional in the AI practices and perspectives they espouse. AI practices around Constructing AI or Analyzing AI are much more present than Communicating about AI. All curricula should balance opportunities for learners to develop all these skills, ensuring that everyone has opportunities to Construct AI and that opportunities to Analyze and Communicate about AI are not overlooked for the sake of constructing AI. In the same vein, All AI perspectives from digital literacy to self-efficacy to expression and activism should be balanced. Emerging priorities around redressing harms of discrimination caused by the tech industry requires shifting from students to be "hackers and makers" to preparing them to be activists and leaders.

3.3.3 Analyzing the UNESCO AI Competency Framework

The UNESCO AI Competency Framework for Teachers and School Students guides teachers on leveraging AI in education and helps students "develop the critical agency, knowledge, skills, attitudes, and values" to navigate AI with agency. Focusing on the student framework, UNESCO describes five main aspects of AI expertise: a human-centered mindset, the ethics of AI, AI foundations, AI skills, and AI for problem-solving. Like the UNESCO Information and Communication Technology (ICT) standards[8], the framework progresses through three levels of competency: knowledge understanding, application, and creation.

8: UNESCO Information and Communication Technology standards, 2018, https://unesdoc.unesco.org/ark:/48223/pf0000265721

Regarding AI content, the framework focuses on general AI knowledge, data and modeling, and the ethics of AI. The "Understand" level of the human-centered mindset, ethics of AI, and AI foundation categories thoroughly cover background knowledge of AI, machine learning, and societal impact. Other AI concepts around machine perception, representation and reasoning, natural communication, and interdisciplinary topics in AI are not covered at all. This makes sense given that the framework focuses on generative AI and education. However, this raises concerns about contextualizing these technologies within the broader practice of AI especially as the field continues to evolve.

The framework takes a relatively broad view of AI practices, particularly constructing and analyzing AI. In the AI foundations, AI skills, and AI for problem-solving sections there are learning outcomes for students learning to conceptualize and implement AI systems in collaboration with others. They also call for students to learn how to evaluate systems for their potential and limitations from sociotechnical perspectives. However, the framework does not include specific design processes, programming tools, or social science analysis tools that students might use to understand AI. Similarly, for communicating about AI, the framework recognizes the importance of students learning to visualize and articulate their understandings of AI systems but does not specify modalities of engagement.

The framework provides moderate coverage of AI perspectives, with the most focus on students' developing awareness and critical consciousness around AI. Particularly in the Apply and Create levels of the progression, the framework strongly specifies ways that students should learn about AI and develop the skills to redesign it to meet their goals. The framework makes a strong bend toward students understanding and advocating for ethical, responsible, and inclusive AI. Specifically, the framework highlights human rights and the importance of students defending their personal and society's communal rights to fairness, human dignity, and agency when interacting with AI. However, there is less focus on the parts of AI perspectives that align with Culturally Responsive Computing, like self-efficacy, action and expression, and community. However, these perspectives are critical if students are to develop positive attitudes toward and identities as technosocial change makers.

Overall, the UNESCO AI Framework is a powerful example of an AI literacy framework that recognizes the inherent sociotechnical nature of AI and strives to have students understand this and take action toward AI that equitably benefits all members of society. As shown in its Impact.AI score card (Table 3.9), the framework prioritizes students developing a critical, yet balanced view of AI. Even so, like many other AI initiatives, it falls short of engaging with students' cultural identities relative to AI or teaching students how they can leverage their knowledge to take action for social justice in society. Impact.AI's explicit attention toward culturally-sustaining pedagogy enables it to directly address barriers to access to participation so that we might empower, include, and encode all students in AI.

AI Perspectives	
E1. Digital Literacy When confronted with new examples of technology, graduates of this course actively leverage their AI knowledge to understand the AI artifact. They actively consider ways that the technology may integrate with their lives and consider how they might redesign the technology to meet their goals.	**E2. Critical Digital Literacy** Graduates of this course think of AI as a force that shapes society and a tool that society shapes in return. They feel that this symbiotic relationship between society and AI can be used to transform society. They also recognize the potential for AI to lead to both benefits and harms. They are aware of their own beliefs about AI and how they relate to others' different perspectives about it.
E3. Digital Citizenship Graduates of this course leverage their AI knowledge to advocate for AI systems that dismantle systems of oppression and improve society. They value diverse perspectives in advocacy for socially beneficial AI and work to recruit and educate others.	**E4. Self-Efficacy** Graduates of this course feel fully empowered (highly motivated and prepared) to design and build new and personally meaningful AI artifacts.
E5. Activism and Expression Graduates of this course recognize their potency as technosocial change agents and actively seek to use technology to expand their understanding and expression of their intersectional identities.	**E6. Community** Graduates of this course feel prepared to operationalize their technosocial agency as part of a community of AI practitioners, activists, and community members. They feel like they belong, can receive help from, can give back to, and help shape this community.

Table 3.8: Criteria for Advanced Coverage of Impact.AI Perspectives. See the full Impact.AI rubric at `https://tinyurl.com/impactai-rubric`

Table 3.9: Impact.AI Score for the UNESCO AI Competency Framework for Students

Item		Score	Comments
		Concepts	
C1.	Background	3.5	Describes AI techniques and applications
C2.	Perception	0	Not covered
C3.	Representation and Reasoning	0	Not covered
C4.	Machine Learning	4	Covers data, algorithms, and models
C5.	Natural Communication	0	Not covered
C6.	Societal Impact	4	Covers critical reflections on AI, human agency, safe and responsible use
C7.	Interdisciplinary Topics	0	Not covered
		Practices	
P1.	Constructing AI	3	Covers ethics by design, AI programming, creating AI projects, co-creation and co-design, problem-scoping
P2.	Analyzing AI	3	Teaches how to evaluate AI models and assess limitations and risks
P3.	Communicating About AI	2	Teaches modeling and visual representations, feedback loops
		Perspectives	
E1.	Digital Literacy	3.5	Promotes awareness of AI in daily life, redesigning AI, safe and responsible use
E2.	Critical Digital Literacy	3	Encourages students to engage in critical reflections on AI, AI and social change
E3.	Digital Citizenship	2.5	Wants students to pursue AI citizenship
E4.	Self-Efficacy	1	Covers growth mindset, resilience, and persistence
E5.	Activism and Expression	2	Wants students to pursue personal fulfillment in the AI era
E6.	Community	0	Not covered

Designing Impact.AI Curricula | 4

This chapter describes the theoretical underpinnings and iterative development of two hands-on middle school AI curricula, "How to Train Your Robot" (HTTYR) and "AI for Wellbeing." HTTYR version 1 was developed and launched in 2019, then entirely redesigned and relaunched in 2020. HTTYR preceded the Impact.AI framework discussed in Chapter 3, but informed much of the framework's goals and content. I began development for AI for Wellbeing in 2022, launching a first pilot in Summer 2023. Although it only underwent one iteration, AI for Wellbeing benefited from the existence of Impact.AI and the lessons learned from developing several K-12 curricula in the intervening years.

This chapter discusses:

- Design principles for K-12 AI curriculum development
- The iterative, participatory design of HTTYR and AI for Wellbeing with teachers
- Teacher's feedback on the curricula
- Implications for designing future curricula given insights from this process

Pieces of this chapter, including the explanation of the design principles, implementation of both iterations of How to Train Your Robot, and design of AI for Wellbeing, build on published papers [110–113]. The explanation of the design principles and technologies used in these curricula also builds on papers I submitted for my 2021 general exams. However, the side-by-side analysis of these curricula with the Impact.AI framework is newly discussed in this work.

4.1 Design Principles for K-12 AI Curricula

In designing K-12 AI curricula created by the MIT Personal Robots Group, we argue for three core design principles: active learning, embedded ethics, and lowering barriers to engagement [112, 114]. These design principles support our goal of creating Impact.AI-aligned curricula that strengthen students' knowledge of AI, develop positive attitudes toward AI, and broaden participation in AI.

[110]: Williams et al. (2021), 'Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum'
[111]: Williams et al. (2024), 'Doodlebot: An Educational Robot for Creativity and AI Literacy'
[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'
[113]: Williams et al. (2023), 'Dr. R.O. Bott Will See You Now: Exploring AI for Wellbeing with Middle School Students'

4.1.1 Active Learning

Active learning is a pedagogical approach where students learn by engaging in activities and then processing new knowledge through reflection [51, 52]. This approach can be effective because it drives students to direct their learning toward knowledge they find compelling.

Active learning begins with hands-on activities, such as plugged or unplugged experiences like online demos, simulations, programming activities, games, and open-ended project design and development. Then, through reflection activities like journaling, surveys, group discussion, and developing and sharing presentations, students engage in spaces where they crystalize and articulate their knowledge.

4.1.2 Embedded Ethics

Embedded ethics refers to teaching technical and ethical knowledge together in lessons rather than separating them [29, 112, 215]. Ethics is a central learning objective included in most K-12 AI frameworks, including the Impact.AI framework [33–35]. The benefit of the practice is that students develop a perspective of technology as inherently sociotechnical, meaning that technologies shape society and vice versa [29, 215]. Practices for embedding ethics include using contextualizing lessons with real-world examples, critiquing AI systems, and teaching students about ethical analysis concepts and practices to boost their understanding of the societal impact of AI [29, 215, 216]. AI ethics is an opportunity to invite students' perspectives and experiences with technology into the classroom to enhance the relevance of the material to students' everyday lives [29].

[29]: Payne (2020), 'Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students'

4.1.3 Lowering Barriers to Access

Lowering barriers to access means designing curricula that mitigate barriers to learning, including engagement, comprehension, and technological access. Our idea of "lowering barriers" builds on Universal Design for Learning [27, 54]. Universal Design for Learning is an approach to education that arose from the disability justice community. It encourages educators to recognize and design around the range of students' interests and learning styles, providing different ways to match lessons with students' strengths. Designing activities that build on students' prior knowledge and offering many opportunities for action are two ways one might lower barriers. Also, ensuring technologies leveraged for learning are affordable, beginner-friendly, and accessible is another important consideration. Different learners and classrooms may differ on what constitutes a barrier to learning. The research team behind HTTYR and AI for Wellbeing primarily worked with Title 1 middle schools in the United States, so we tailored our design choices to their needs. Some assumptions we designed around were that lessons would be more effective in the classroom rather than in extracurricular spaces, that we would need to provide free teacher training and technological resources to classrooms, and that students would be new to programming. Table 4.1 shows examples of how each of the three design principles took shape in our two curricula.

Design Principle	How to Train Your Robot	AI for Wellbeing
Active Learning	In the introduction to"Image Classification" activity, students experimented with the Quick Draw game and dataset to see how different features of images contribute to how algorithms learn to distinguish classes.	On the first day, students engaged in a class discussion where they identified everyday examples of AI technologies and their benefits and harms. Examples were collected on a shared worksheet to return to throughout the week.
Embedded Ethics	All students completed and presented ethical matrices for their final projects. In the matrices, students identified the key stakeholders and potential good and bad consequences of deploying their projects.	In the"Intro to Design Justice" activity, students explore the possibilities of different AI technologies, like smart assistants and virtual doctors, from the lens of different stakeholders.
Lowering Barriers to Access	Since there is no programming prerequisite for the curriculum, the first programming activity uses a set of mini challenges to get students up to speed on the important robotics and programming concepts they will need later in the class.	A virtual Jibo is built into the RAISE AI Playground to accommodate users who do not have access to physical Jibo robots. Care was taken to make sure that the programming blocks for controlling physical Jibo are compatible with the virtual version.

Table 4.1: Design Principles for K-12 AI Curricula applied to How to Train Your Robot and AI for Wellbeing

4.2 Design Methodology

As discussed in the Background (Chapter 2, Section 1) and Impact.AI Literature Review (Chapter 3, Section 2), research on K-12 AI curricula in classrooms with in-service teachers is currently lacking. Many researchers have pointed out that research-practitioner partnerships are critical for creating opportunities for all students to learn about AI [173, 208, 217]. Thus, when our research team undertook to design K-12 AI curricula, we partnered with in-service middle school teachers. We leveraged an iterative design research methodology in collaboration with K-12 educators to develop and evaluate our curricula. In our approach, the cyclical process of development, participation of educators, and focus on usability all played a major role in the method and pace of our design [218]. First, we developed the curriculum, often in collaboration or conversation with educators who set our design focus. Then, we would train educators to use the curriculum and have them run it with students. We used surveys and interviews to collect teachers' feedback then leveraged the feedback to develop the next prototype.

Through this design methodology we ought to answer the following research question: What are key design considerations for incorporating AI curricula in middle school classrooms?

4.2.1 Overview of Design Context

The original impetus for running these curricula was to fulfill a curriculum request from I2 Learning for their 2019 Massachusetts STEM week. STEM Week is a STEM outreach program, organized by I2 Learning, where middle school classrooms in Massachusetts do hands-on, intensive STEM courses for an entire week. I2 Learning allowed us to create their first AI STEM week curriculum.

A STEM week curriculum consists of 30 hours of hands-on activities culminating in a final design project. To prepare teachers for STEM week, teachers complete a two-day intensive training and receive a comprehensive educator guide, worksheets, and a slide deck. Workshops were run every single day of the week, Monday to Friday, with the course culminating on Friday with a project showcase.

We originally designed How To Train Your Robot in this context, launching with them in October 2019. In 2020, we revised HTTYR. Then, due to the health caused by the COVID-19 pandemic, we modified this second version of HTTYR from an in-person format into synchronous online summer workshops where we invited students and teachers to join us for training and to learn about AI. Williams et al. (2022) describe how we modified lessons as we moved to the online learning format [112]. We continued this format in 2023 with a new curriculum, AI for Wellbeing.

[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'

4.2.2 Partnerships with Teachers

Across all three curricula, teachers were deeply embedded in the design and development of the curricula. The original idea for designing the curricula came from the I2 Learning group, and they brought to us their ideas for revamping a curriculum previously designed by Blakeley H. Payne (2020), a former member of the Personal Robots Group, to match their existing format [29]. The teachers we worked with throughout each iteration of the curricula were simultaneous trainees earning professional development credit for their work and participants in the evaluation and feedback of the curriculum. We also worked closely with an I2 Learning Curriculum Developer, Sam Forman, an educator with many years of teaching experience, to revise How to Train Your Robot for its second version.

[29]: Payne (2020), 'Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students'

4.3 Iterative Design of K-12 AI Curricula

Each curriculum consists of activities, delivered to teachers in the form of an educator guide and accompanying hardware and software platforms. Figure 4.1 shows the first version of How to Train Your Robot was developed and tested in 2019. Then, the author and a curriculum developer for I2 Learning revised the curriculum for STEM Week 2020. The evaluative study of the 2020 version of How to Train Your Robot curriculum was conducted online with the Amazon Future Engineers (AFE) program. Between 2020 and 2023 we used How to Train Your Robot as a foundation for different forms of K-12 curricula, a topic which we will further describe in the Discussion (Section 5) of this chapter. Finally, we came back to developing weeklong AI curricula and we created AI for Wellbeing curriculum for STEAM Ahead's STEAM Academy.

In this chapter, we will discuss each of the three AI curriculum prototypes we developed. We will cover the design of the activities and the hardware and software platforms we developed for the curricula; then we will share feedback from the teachers we trained and engaged with to do the study.

4.3.1 First Prototype: How to Train Your Robot v1.0

The design and results of this curriculum were originally published in full in Williams, Kaputsos, and Breazeal (2021) [110]. We are restating them here for comparison.

[110]: Williams et al. (2021), 'Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum'

Activities

The original activities developed for How to Train Your Robot built heavily on the Middle School AI + Ethics Curriculum developed by Blakeley H. Payne (Payne 2020) [29]. In this curriculum, Payne embedded ethical lessons about AI with technical ones, raising students' awareness of the ethical implications of AI technologies and showing them how to navigate the sociotechnical nature of all technology. Similarly, every How to Train Your Robot curriculum session incorporated technical and ethical concepts.

[29]: Payne (2020), 'Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students'

In the technical lessons, students learned about general definitions in AI, text and image classification, and training supervised machine learning models. In the ethics modules, students learned about algorithmic bias, AI's positive and negative impacts, product design trade-offs, and how to conduct stakeholder analysis using ethical matrices. Each session focused on real-world examples and hands-on opportunities to see AI algorithms. At the end of the week, students completed a final project, applying their technical and ethical skills to projects they cared about.

Session	Activity	Learning goals
Session 1	AI or Not	Define AI, reason about what makes something AI or Not
	Ethical Dilemmas Video	Employ decision-making strategies to reason through ethical dilemmas
	Intro to ScratchX	Use block-based programming to complete mini-projects with robots
Session 2	PB&J Recipe Algorithm	Define algorithms and design an algorithm
	Algorithmic Bias Video	Discuss the implications of algorithmic bias and what can be done to mitigate it
	Image Classification	Curate datasets and use them to train image classifiers, understand neural networks
	ML4Kids + ScratchX	Use ML4Kids to create program robots to play a simple card game
Session 3	Ethical Matrices	Use stakeholder analysis to unpack the implications of technology design
	Self-Driving Robots	Learn about closed-loop algorithms to have robots autonomously navigate a path
	Final Project Research	Do research to generate ideas for final projects
Session 4	Final Project Planning	Use project planners and ethical matrices to develop final project ideas
	Final Project Work Time	Employ time management to work on an open-ended project
	Final Project Peer Review	Give and receive peer feedback
Session 5	Final Project Showcase	Complete a final project that applies what students learned during the course Create a presentation to explain final project to a general audience

Table 4.2: An overview of the activities and learning goals in How to Train Your Robot v1.0

Table 4.3: Impact.AI Score for HTTYR v1.0

Item		Score	Comments
Concepts			
C1.	Background	3.5	AI definition, examples of AI, AI benefits and harms
C2.	Perception	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C3.	Representation and Reasoning	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C4.	Machine Learning	3	Covered supervised classification algorithms
C5.	Natural Communication	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C6.	Societal Impact	3	Described AI benefits and harms of AI based on different stakeholders, algorithmic bias
C7.	Interdisciplinary Topics	0	Not covered
Practices			
P1.	Constructing AI	2.5	Implementing AI projects in ScratchX
P2.	Analyzing AI	3	Teaches how to evaluate AI models and analyze AI systems with stakeholder analysis
P3.	Communicating About AI	2	Students present final AI projects to general audience
Perspectives			
E1.	Digital Literacy	2.5	Promotes awareness of everyday examples of AI, distinguishing AI or not
E2.	Critical Digital Literacy	3	Promotes awareness of AI's benefits and harms
E3.	Digital Citizenship	1	Encourages students to engage with AI for social good
E4.	Self-Efficacy	3.5	Hands-on practice building AI projects with mentors
E5.	Activism and Expression	2.5	Using AI for social good
E6.	Community	1	Engage in class community, share final projects with families and friends

Hardware and Software

The first version of How to Train Your Robot used an extendable version of Scratch called ScratchX. ScratchX is no longer functional since it relied on Adobe Flash Player, which was deprecated in 2020. We originally chose this software because it was a free, browser-based, block-based programming language that can be connected to extensions. The extension we used for training machine learning models was the Machine Learning for Kids (ML4Kids) website. Teachers had to register students for a trial developer account through IBM Watson's cloud API to use this tool. Then, students can use a kid-friendly interface to image, text, audio, and number classification models and control them with Scratch blocks.

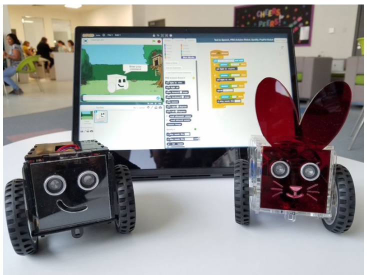

Figure 4.1: HTTYR v1.0 used arduino-based Gizmo robots and the ScratchX programming website to give students hands-on experiences programming machine learning projects

Although the focus of the curriculum was not robotics, we included a robot because physical manipulatives have been known to increase student engagement and help them connect with abstract ideas [48, 219–222]. We built a $50 Arduino robot, called Gizmo, based on an open-source project called PopPet created by Jaidyn Edwards. We could control Gizmo using a ScratchX extension based on an Arduino-controlling extension created by Kreg Hanning. The robot included line sensors, ultrasonic distance sensors, motors, and RGB LEDs.

4.3.2 Second Prototype: How to Train Your Robot v2.0

Like the first version of How to Train Your Robot, the second version is also described in Williams, Kaputsos, and Breazeal (2021) [110]. This section is a summary of what is found in that paper.

[110]: Williams et al. (2021), 'Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum'

Activities

In the second version of How to Train Your Robot, we worked closely with Sam Forman, a middle school teacher and former curriculum developer at I2 Learning. In the redesign, we prioritized reinforcing students' understanding of fundamental AI concepts, like understanding the definition of AI, accommodating students and teachers new to programming, tying in real-world relevance, and using more robust technological tools. We changed activities on the second day to give students additional opportunities to explore algorithms by designing an Alexa Pizza Delivery app. For the third day, we added lessons on text classification with a hands-on programming activity.

Table 4.4: An overview of the activities and learning goals in How to Train Your Robot v2.0

Session	Activity	Learning goals
Session 1	AI or Not	Define AI, reason about what makes something AI or Not
	Ethical Dilemmas Video	Employ decision-making strategies to reason through ethical dilemmas
	Intro to the AI Playground	Use block-based programming to complete mini-projects with robots
Session 2	Alexa Pizza Delivery App	Define algorithms and design an algorithm
	Image Classification + Algorithmic Bias	Curate datasets and use them to train image classifiers, discuss the implications of algorithmic bias
	Teachable Machine + the AI Playground	Build custom image classification models to program robots in the AI Playground
Session 3	Ethical Matrices	Use stakeholder analysis to unpack the implications of technology design
	Exploring Word Analogies	Understand word vectors and how they encode language, discuss bias in large datasets
	Text Classification + the AI Playground	Build custom text classification models to program robots in the AI Playground
Session 4	Final Project Brainstorming	Conduct research to generate ideas for final projects
	Final Project Planning	Use project planners and ethical matrices to develop final project ideas
	Final Project Work Time	Employ time management to work on an open-ended project
Session 5	Final Project Showcase	Complete a final project that applies what students learned during the course
		Create a presentation to explain final project to a general audience

Table 4.5: Impact.AI Score for HTTYR v2.0

Item		Score	Comments
Concepts			
C1.	Background	3.5	AI definition, examples of AI, AI benefits and harms
C2.	Perception	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C3.	Representation and Reasoning	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C4.	Machine Learning	3	Covered supervised classification algorithms
C5.	Natural Communication	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C6.	Societal Impact	3	Described benefits and harms of AI based on different stakeholders, humans' role in AI, algorithmic bias, critical discussions of AI
C7.	Interdisciplinary Topics	0	Not covered
Practices			
P1.	Constructing AI	3.5	Problem scoping, design thinking, implementing AI projects in AI Playground
P2.	Analyzing AI	3	Teaches how to evaluate AI models and analyze AI systems with stakeholder analysis
P3.	Communicating About AI	2	Students present final AI projects to general audience
Perspectives			
E1.	Digital Literacy	2.5	Promotes awareness of everyday examples of AI, distinguishing AI or not
E2.	Critical Digital Literacy	3	Promotes awareness of AI's benefits and harms, exposure to different viewpoints about AI
E3.	Digital Citizenship	1	Encourages students to engage with AI for social good
E4.	Self-Efficacy	3.5	Hands-on practice building AI projects with mentors
E5.	Activism and Expression	2.5	Using AI to make a difference, working on personally meaningful problems
E6.	Community	1	Engage in community as class, share with others

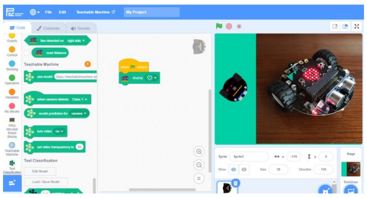

Figure 4.2: HTTYR v2.0 used micro:bit robots, either Tinybit or Cutebot, and the RAISE AI Playground to give students hands-on experiences programming machine learning projects

[69]: Jordan et al. (2021), 'PoseBlocks: A Toolkit for Creating (and Dancing) with AI'
[71]: Reddy et al. (2021), 'Text Classification for AI Education.'

Hardware and Software

For How to Train Your Robot 2.0, we moved to the RAISE AI Playground, our fork of the open-source Scratch 3.0 repository. In this repository, we built-in extensions for Google's Teachable Machine, a new text classification extension, and extensions that control micro:bit robots. Teachable Machine is a tool from Google that makes it easy to train and run inference on image recognition on a browser without a service account [65]. The Teachable Machine extension developed by Jordan et al. (2021) allows users to load and control Teachable Machine models from block-based programs [69]. We also developed a new extension to train supervised machine learning models for text classification. The interface, published in Reddy, Williams, and Breazeal (2021) [71], allows students to train and run models directly in the RAISE AI Playground interface. Moving to this interface allowed students to create and use unlimited models in a streamlined interface where all extensions were in one place. Since How to Train Your Robot 2.0, we have continued developing new extensions in the RAISE AI Playground to teach new AI subjects.

We also moved from hand-built robots to commercial micro:bit robots from Yahboom and Elecfreaks that cost about $40 each. Like the Gizmo robots, the robots had line sensors, ultrasonic distance sensors, motors, and RGB LEDs. The robots also included a piezo buzzer, two push buttons, and a 25-LED display. Unlike Gizmo, these robots were reliable and used a newer Bluetooth interface that was much easier to connect to students' computers.

4.3.3 Third Prototype: AI for Wellbeing 2023

AI for Wellbeing was submitted for publication at the Educational Advances in Artificial Intelligence (EAAI) conference 2024. That publication described the curriculum design and preliminary analysis of teachers' feedback on the curriculum. In this section, we build on that prior publication with further analysis of teachers' responses.

Activities

After several years developing the Impact.AI framework, the Day of AI outreach initiative, and a semester-long co-designed AI curriculum, we returned to the weeklong AI curriculum format in 2023. With AI for Wellbeing, we partnered with STEAM Ahead to do a one-week, online summer camp. We prioritized developing a new curriculum, updated to cover emerging AI topics and more intentionally designed to reflect a culturally responsive pedagogy approach. Except for the final project, most of the activities in the curriculum were heavily edited or newly developed.

The activities in Session 1 included having students create a basic definition of AI, learn about everyday examples of AI, and understand AI's potential benefits and harms (C1. Background in AI, C6. Big Idea #5 Societal Impact). We discussed everyday benefits and harms of AI so that students would better recognize examples of AI and develop a critical awareness of their limitations (E1. Digital Literacy, E2. Critical Digital

Literacy). The lessons in Units 2 and 3 covered fundamental concepts in supervised machine learning (C4. Big Idea #3 Machine Learning) and conversational agents (C5. Big Idea #4 Natural Communication). In addition to understanding the technical implementations of these technologies, we also discussed ethical issues related to chatbots, like algorithmic bias and privacy concerns.

For practices, we wanted students to be able to design and implement conversational agents (P1. Construction AI) and use the Design Justice framework to analyze AI systems (P2. Analyzing AI). Students practiced their computational thinking skills daily by programming Jibo, the social robot. Students also implemented rule-based and neural-network-based machine learning algorithms to turn their Jibo robots into "intelligent" social companions. Design Justice is a methodology that puts marginalized communities at the forefront of design to challenge the inherent inequality in the design process [28]. Students learned a simplified version of the framework, based on a lesson created by Blakeley H. Payne (2021) where students investigate AI artifacts by considering "Who participated?" "Who benefited?" and "Why was harmed?" [223].

[28]: Costanza-Chock (2020), *Design justice: Community-led practices to build the worlds we need*

[223]: Payne (2021), 'Bringing design justice to the classroom'

Students' final project work allowed them to build confidence in their skills (E4. Self-efficacy) and practice using AI to express themselves or take action in their community (E5. Action and Expression). In students' final projects, they apply their new technical and critical analysis skills to projects they are personally motivated to complete. The workshop concluded with students sharing one another's projects (P3. Communicating about AI).

Table 4.6: An overview of the activities and learning goals in AI for Wellbeing

Session	Activity	Learning goals
Session 1	AI or Not	Define AI, reason about what makes something AI or Not, learn about AI benefits and harms
	Intro to the AI Playground	Use block-based programming to control Jibo robots
Session 2	Intro to Design Justice	Use the design justice framework to redesign classroom policies
	Intro to Text Classification	Assemble datasets to train and test text classification models
	Text Classification + the AI Playground	Build custom text classification models to program Jibo in the AI Playground
Session 3	Chatbot Design	Understand different components of chatbots, explore interface design
	Final Project Brainstorming	Use brainstorming tool to generate ideas for final projects, give and receive peer feedback
	Chatbots + the AI Playground	Create dialogue flows to interact with Jibo in the AI Playground
Session 4	Final Project Planning	Use project planners and ethical matrices to develop final project ideas
	Final Project Work Time	Employ time management to work on an open-ended project
Session 5	Final Project Showcase	Complete a final project that applies what students learned during the course Create a presentation to explain final project to a general audience

Table 4.7: Impact.AI Score for AI for Wellbeing

Item		Score	Comments
Concepts			
C1.	Background	3.5	AI definition, examples of AI, AI benefits and harms
C2.	Perception	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C3.	Representation and Reasoning	1	Briefly covered all 5 Big Ideas in the Intro to AI activity
C4.	Machine Learning	3	Covered supervised classification algorithms
C5.	Natural Communication	3	Covered chatbots, natural language processing, human-computer interaction
C6.	Societal Impact	3	Described AI benefits and harms of AI based on different stakeholders, human role in AI, algorithmic bias, critical discussions of AI
C7.	Interdisciplinary Topics	0	Not covered
Practices			
P1.	Constructing AI	3.5	Problem scoping, design thinking, implementing AI projects in the AI Playground
P2.	Analyzing AI	3	Teaches how to evaluate AI models, analyze AI systems with the design justice framework
P3.	Communicating About AI	2	Students present final AI projects to general audience
Perspectives			
E1.	Digital Literacy	3	Promotes awareness of everyday examples of AI, distinguishing AI or not
E2.	Critical Digital Literacy	3	Promotes awareness of AI's benefits and harms, exposure to different viewpoints about AI
E3.	Digital Citizenship	3	Encourages students to engage with AI for social good, welcoming more voices into AI design process
E4.	Self-Efficacy	3.5	Hands-on practice building AI projects with mentors
E5.	Activism and Expression	3	Using AI to address societal inequality, working on personally meaningful problems
E6.	Community	1	Engage in class community

Figure 4.3: AI for Wellbeing used the Jibo social robot and the RAISE AI Playground to give students hands-on experiences programming chatbots

Hardware and Software

We continued to use the RAISE AI Playground in the AI for Well-being curriculum. We used the text classification extension, updated to include additional natural language processing blocks and a newly developed Jibo robot extension. We also integrated two previously developed AI scaffolding tools, Sparki and LevelUp, that students could use to get feedback on technical and ethical aspects of their projects as they worked. We will further discuss these two tools in Chapter 6.

Students used the Jibo robot, an autonomous social robot designed for the home. In its autonomous mode, Jibo's vibrant personality entertains users with games, plays music and dances, and answers questions based on information from the Internet. In its programmable mode, users can use the AI Playground to make Jibo talk, detect voice input, dance, change Jibo's LED light, and show various icons on its screen. Students participating in the work lived with the Jibo social robot for about two weeks, allowing them to gain experience with the robot they were programming. Although most students could use the Jibo robot, a few used a virtual version of Jibo built into the AI Playground.

4.4 Teacher's Feedback on Curriculum Prototypes

4.4.1 Results from HTTYR v1.0

HTTYR v1.0 was run much differently from HTTYR v2.0 and AI for Wellbeing. This analysis of teachers' feedback first looks at teachers' feedback from the former, and then compares the other two prototypes side-by-side.

Participants

In this first iteration, we recruited seven teachers to run the curriculum as part of Massachusetts STEM Week. Of those seven, three teachers from two schools consented to be research participants and complete surveys before and after using the curriculum.

Procedures

Before STEM Week, teachers underwent a two-day, intensive training program with the researchers who designed the curriculum. The goal was for them to try all the activities they would do in the class. One issue was that the Arduino robots used in the class were unavailable during the training program, so teachers had to use a substitute robot.

During STEM Week, teachers spent about 25 hours across five days going through the curriculum in their classrooms. The researchers who designed the curriculum were available to teachers to answer questions and resolve technical issues. The most issues came up on the first day

Participant	School	Subjects Taught	AI Experience
P1 (Male)	School 1 - Title 1 school in an urban area	Math	Comfortable with technology, never taught computer science before
P2 (Female)	School 2 - Small school in a rural district	Many subjects	Comfortable with programming
P3 (Female)	School 2 - Small school in a rural district	Science and Math	Less comfortable with technology

Table 4.8: Description of teachers who participated in the HTTYR v1.0 curriculum study

when teachers had to connect robots to the computers, and on the fourth day, the software students were using had a temporary outage.

After the curriculum, we interviewed teachers using 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 were some lowlights 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?

After the interviews, two researchers reviewed teachers' responses to identify common themes in their responses.

Results

The three major themes we identified in the teachers' responses centered around student engagement and using robots in the classroom.

Student engagement. Teachers reported that their students were very engaged in learning about AI, alluding to students being "digital natives" with prior experiences and knowledge of AI (cite Mimi Ito).

> "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." (P2)

However, as seen in the literature about STEM identity, some students felt more prepared to engage in AI than others.

> "Some students sat back and let others do most of the work." (P2) The embedded ethics approach of our curriculum was meant to ensure that all students got something out of the activities. However, there were mixed responses about how

> engaging these activities were. "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)

Robots in the classroom. We saw that the hardware was an exciting tool for teachers to use in the classroom, but it did cause extra difficulty for teachers.

> "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)

Conclusions. From the teachers' feedback, we identified three key ways to improve the curriculum for the subsequent year. First, we wanted to build on the activities that worked well, like AI or Not, and improve the activities that caused a lot of difficulty, like programming the robots. The ScratchX interface with Machine Learning for Kids required too much setup, and the hand-built Arduino robots needed to be more reliable for classroom use. By redesigning the activities and streamlining the use of technology, we made it easier for teachers to navigate the activities successfully.

4.4.2 Results from HTTYR v2.0 and AI for Wellbeing

Participants

For HTTYR v2.0 we recruited seven teachers, and for AI for Wellbeing we recruited five teachers. All twelve teachers consented to participate in the research portion of the workshop and complete surveys before and after the workshop, as well as interviews. Teachers were compensated $500 for their participation in the workshops, regardless of whether they consented to participate in the research portion.

We specifically recruited in-service teachers in Title 1 schools in the United States. Title I is a designation given to schools that receive financial assistance from the federal and local government because their student body consists of high percentages of children from low-income families. In HTTYR v2.0, teachers came from suburban, urban, and rural schools. One teacher directed a homeschool network that specifically served students with special needs. Every teacher we recruited taught a STEM subject, most of whom taught Computer Science. Only two teachers had prior experience teaching AI before, but all expressed interest in bringing AI to their classrooms to serve their students better.

The teachers recruited for AI for Wellbeing were about half from Title 1 schools and half from alternative education programs, including STEAM-focused, diversity afterschool programs. At least one teacher

Participant	School	Subjects Taught	AI Experience
P1 (Female)	Title 1 public school, suburban area	Science, Aerospace, Robotics	No prior experience teaching AI
P2 (Female)	Title 1 public school, suburban area	Computer Science, STEM	No prior experience teaching AI
P3 (Male)	Title 1 public school, suburban area	Computer Science	Has taught AI before
P4 (Female)	Title 1 public school, urban area	Computer Science	Taught AI to high schoolers before
P5 (Female)	STEM magnet school, rural area	Technology, Engineering	No prior experience teaching AI
P6 (Female)	Directs home-school network for students with special needs	Computer Science, Computer Applications, STEM	No prior experience teaching AI
P7 (Female)	Title 1 school, all female	Computer Science, Engineering	No prior experience teaching AI

Table 4.9: Description of teachers who participated in the HTTYR v2.0 curriculum study

had experience working with students with special needs, specifically students with Autism Spectrum Disorder. At least one teacher had experience working with students from non-English speaking homes. Most teachers taught Computer Science or Technology, except for one teacher who focused on art but used technology to explore art in their classroom.

Procedures

For HTTYR v2.0 and AI for Wellbeing, we ran the evaluative students as online, summer workshops for students. The format of these workshops were 2 - 2.5 hour sessions every day for a week, Monday to Friday. Rather than doing a separate teacher training before the workshop, researchers ran through the next day's activities for daily training after class. During class time, researchers delivered brief introductions to the topics, organized transitions between activities, and supported teachers. Pairs of teachers worked together to lead students through activities in breakout rooms. In total, we spent about 12 hours doing the curriculum with students and 5 hours in teacher training, significantly less time than the first version of the curriculum. After completing the training, teachers were invited to continue the curriculum in their classrooms the following year in their classrooms.

After each day, we asked teachers:

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

Table 4.10: Description of teachers who participated in the AI for Wellbeing curriculum study

Participant	School	Subjects Taught	AI Experience
P8 (Female)	Alternative education program	Art, Storytelling and Design	Explored AI tools in classroom before
P9 (Female)	Title 1 public school, urban area	Computer Science, Technology	Taught AI before, not to middle schoolers
P10 (Male)	Title 1 public school, suburban area	Computer Science	Has taught AI before
P11 (Male)	Alternative education program	Computer Science	Has taught AI before
P12 (Female)	Title 1 school, suburban area	Digital Information Technology	No prior experience teaching AI

After the curriculum, we interviewed teachers using the following questions:

1. Did the material in this course change your opinion about AI or teaching AI to students?
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 would bring to this course?
5. Which parts of this course would you bring to your classroom, and which parts would you leave behind?

After conducting interviews for HTTYR v2.0, two researchers engaged in a qualitative analysis of teachers' responses using an open coding approach. The process involved each researcher independently reviewing interview transcripts and assigning one of five codes to each quote. These codes, elaborated in Table 4.11, were selected based on topics that came up in the HTTYR v1.0 workshop and insights from Vazhayil et al. (2019) [217], who conducted a teacher training program with middle school teachers. Any discrepancies between researchers were resolved through collaborative discussion.

Following the coding process, the researchers collaborated to distill overarching themes that captured teachers' perspectives. These themes were: the effectiveness of the curriculum design, the effectiveness of the technology, strategies for supporting students, and recommendations for improvement and future curricula. Once researchers reached consensus on these themes, they clustered quotes with similar ideas to discern the prevalence of different opinions across the participants. The findings from the HTTYR v2.0 workshop were published in Williams et al. (2021) [110].

[110]: Williams et al. (2021), 'Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum'

In 2023, after the AI for Wellbeing workshop, one researcher returned to the themes derived from the prior study to analyze the responses of the new cohort of teachers. While many themes from the HTTYR v2.0 study persisted, a new theme emerged concerning teachers' motivation to teach AI, while the theme regarding supporting students became less

Table 4.11: Description of codes used to organize quotes from teacher interviews after the HTTYR v2.0 workshop

Code		Examples
technology	refers to hardware or software tools used in the workshop	"I liked using Scratch" "Google Classroom was confusing" "I think the robots were the best part"
strategy	refers to a pedagogical strategy that was used in the classroom.	"Splitting up students into smaller groups allowed them to participate more" "Getting students to apply their knowledge to projects took a lot of effort"
engagement	refers to students being interested or engaged disinterested or not paying attention	"I think students were most excited when they got to redesign the technology themselves" "Students were quiet and did not talk"
insight	refers to teachers / students learning understanding or gaining knowledge or perspectives	"They had never thought of it that way" "I don't think my explanation made a whole lot of sense" "They learned how to debug their code" "My students learned to be persistent and to keep trying"
training	refers to some material or resource that was used to prepare teachers in the class	"I couldn't answer their questions because we didn't talk about it as teachers" "The educator guide made it clearer"

relevant. I was interested in pursuing the new theme about teachers' motivation to teach AI because of the work's new focus on identity and inclusion. To update the HTTYR v2.0 analysis, one researcher revisited teacher's quotes from the prior study to identify quotes that aligned with the new theme. I omitted the theme related to supporting students in these results because the HTTYR v2.0 workshop occurred in 2020 and this theme primarily addressed new challenges of adapting to virtual classrooms, however this no longer felt relevant in 2023.

The results section below summarize teachers' statements about their motivations for teaching AI, students' learning and engagement during the workshops, the effectiveness of the technology, and teachers' interest in expanding the curriculum.

Results

Teacher's Motivation to Teach AI

All teachers recruited for HTTYR v2.0 and AI for Wellbeing stated that they were comfortable with technology, and many were teaching STEAM courses. In HTTYR v2.0, only two of 7 teachers had prior experience teaching AI in the classroom. That did not change much by 2023 with the teachers who participated in AI for Wellbeing, despite the increased availability of free tools and courses for teaching and learning about AI.

Even so, many of the HTTYR v2.0 teachers (5 out of 7) were already creating advanced computing models for the classes and extracurricular

STEM activities they were engaging in with students. Teachers expressed their desire for their students to get exposure to emerging areas of technology or further pursue their existing interests in computing.

> “Our students enjoyed learning about artificial intelligence during this past year’s Hour of Code. Participating in this workshop would allow them to continue learning about this rapidly changing field of study." (P02)
> “I have done a lot of personal work with AI and I think it is definitely something that the students need to learn as more and more of society transitions into an AI-centric world. This will help them be competitive" (P03)
> “Students are eager for hands on experiences and learning new things that could possibly shape their future" (P04)

Some teachers explicitly expressed that their students being women, from minoritized racial groups, or having disabilities motivated their interests in students getting further exposure.

> “My students and I represent populations that have been historically ignored and marginalized in the STEM industry. This will be a great opportunity for them to experience this program." (P06)
> “My motivation to teach AI is grounded in promoting equity and inclusivity. I am driven to bridge the educational gap, especially for historically underrepresented communities." (P11)

Many teachers came to the course with some alignment to the goals of Impact.AI. They wanted students to learn about AI’s practical applications in daily life, practice creativity and problem-solving, and understand technology’s societal impact. Teachers who participated in the AI for Wellbeing curriculum were especially aligned in their thinking. This was not surprising since many of them had used other resources from MIT RAISE before, particularly the Day of AI curriculum.

> “I hope that students will recognize biases in computer programs so as they learn more they will be better creators in tech than we (adults) are currently." (P9)
> “I want them to foster critical thinking and problem-solving abilities to approach AI challenges creatively. Understanding the societal implications of AI, particularly within marginalized communities, is crucial." (P11)
> “[My goal is] to have my students gain experience and learn more about AI and how they can use it in their daily lives or have them explore with better ways to help others." (P12)

Student Learning and Engagement

After both curricula, teachers expressed that the curriculum effectively engaged their students. In How to Train Your Robot, teachers were impressed to see their students continue working on projects 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)

In AI for Wellbeing, students and teachers had not met previously, so there were fewer opportunities for class participants to come together outside of class. Even so, the majority of teachers selected "Agree" or "Strongly Agree" to the statement "My students were engaged in the AI workshop."

Teachers from both curricula consistently mentioned that the programming activities were most engaging for their students (see Figure 4.5). In HTTYR v2.0, they felt that students having the opportunity to participate in the full AI development process, from training to testing and then deploying models in their projects, was critical to their engagement.

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

In AI for Wellbeing, playing with Jibo seemed to engage students most.

> "I think the students got a kick out of all of [Jibo's] capabilities." (P09)

An issue that came up in HTTYR v1.0 which was addressed in HTTYR v2.0 was the mixed reviews on the ethics activities. Some teachers believed they were invaluable, others felt they were awkward and confusing for students. In the broader context of teachers' backgrounds, it was clear that teachers' readiness to engage in candid conversations about race influenced their perceptions about the ethics activities. For instance, a general education teacher praised the algorithmic bias discussion as one of their favorites, and connected it to a prior lesson on racial justice from another subject. Conversely, a mathematics teacher, who admitted they were not accustomed to discussing race in their class, said that students were hesitant to open up and that the experience was awkward.

Recognizing that classroom discussions about discrimination can be intimidating, HTTYR v2.0 took steps to better equip students to engage in critical conversations. The revisions included multiple options for exploring the presence and impact of discrimination through code samples, videos, and articles. In addition to classroom discussion, we also suggested activities like journaling to offer different pathways toward honest, yet safe ways to address challenging topics. The goal was to accommodate teachers at different levels of readiness to foster authentic engagement with the topic. In addition to these revisions, HTTYR v2.0 teacher participants also benefited from support from the researchers and their peers during the algorithmic bias activities. Overall, teachers' feedback on the second iteration of ethics activities was uniformly positive.

> “Having them be able to see on paper the different people involved, that’s going to be valuable. . . It helps them to 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)

Similarly, in AI for Wellbeing, teachers seemed to value the ethics lessons. The Intro to Design Justice activity was a particularly effective ethics activity because it was designed to have students practice their design thinking skills on a highly relevant and important topic, their classroom rules.

> “[The most engaging moment] is the discussion of deciding the class policy. They enjoy the feeling of being in charge." (P08)

However, we maintain that ethical reasoning skills are difficult for both teachers and students to master, especially when they are new to the practice. It is important to support and scaffold students in this space and for technical skills.

> “I don’t think kids can easily identify who would be harmed by AI or what the potential harms might be." (Anonymous, AI for Wellbeing)

The apparent perception among teachers that the technology activities were more engaging than the ethics ones raises an important concern about embedding ethics in AI curricula. In STEM courses, ethics topics are often deemed tangential to the core curriculum and are at risk of being sidelined. However, this work advocates for the persistent inclusions of ethics activities, and particularly those that address systemic discrimination. While programming a physical robot may enhance students’ self-efficacy and empowerment, engaging in conversations that openly addresses inequality and confronts stereotypes contributes to increased self-efficacy, empowerment, and a sense of belonging. Ethics activities offer a more profound and enduring form of engagement than programming activities. They help students foster critical consciousness, shaping how they perceive themselves in relation to technology and society.

Effectiveness of the Technology

Technology was always a double-edged sword in every iteration of the curricula, creating moments of engagement and frustration. In HTTYR v2.0, teachers enjoyed that the programming was built into the Scratch interface since many students were familiar with it. The added functionality of the RAISE AI Playground was seen as a significant improvement to the official Scratch website.

> “Customized Scratch allowed me to [understand AI] and introduce it to 5th graders." (P06)
> “Making it Chromebook accessible was perfect for this time"

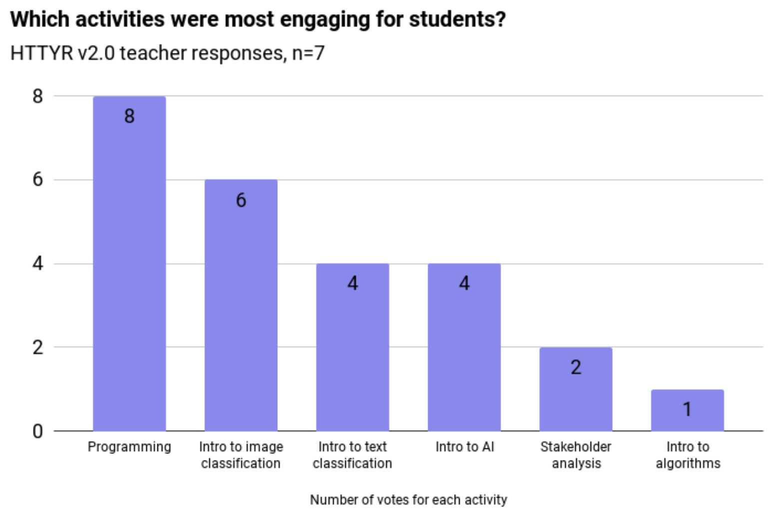

Figure 4.4: Teachers reported which activities in the HTTYR v2.0 curriculum they thought were most engaging and effective for students

> "Scratch has been in the school for a long time." (P07)

However, not all students had experience with Scratch and every classroom had students with mixed programming abilities. In the AI for Wellbeing course, limited class time created an issue where diving deeper into the AI activities would come at the expense of the students building a stronger foundation with the programming.

> "We all had trouble with the coding and text classification." (P09)
> "[The most difficult part] was setting up the [text classification] classes, getting the Scratch program to respond appropriately after inputting some form of hello or goodbye." (P10)
> "[One thing I would add to this curriculum is] maybe more hands-on time for skill building. . . also some questions were on operating Playground itself. (How to change language, how to save files, how to connect the robot name to the code etc)" (Anonymous, AI for Wellbeing)
> "[One thing I would add to this curriculum is] vocabulary and an explanation of what each thing does in [the AI Playground]." (Anonymous, AI for Wellbeing)

Handling this issue is a subjective and class-dependent. In HTTYR v2.0, teachers and students could spend time outside of class to review and practice topics before coming to class the next day. In AI for Wellbeing, we slowed down the class and removed some activities from Day 3 to create more space for students to grasp text classification and programming.

Given the online context of these workshops, we also considered breaking students up into groups based on their ability, a strategy we had used in prior workshops that seemed effective (Williams et al., 2022). However, we gave students the power to make that decision in the Intro to Design Justice activity, and they chose to stay in the same breakout groups for the week to become more familiar with a smaller group and learn to support one another. There are a variety of approaches that educators

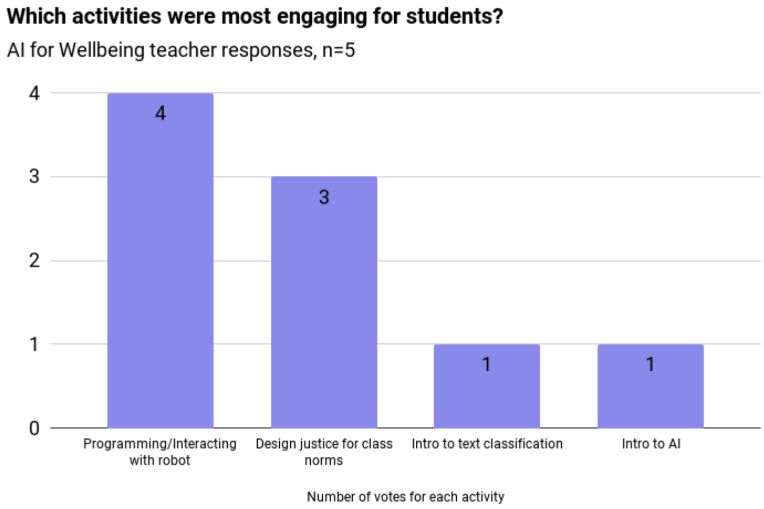

Figure 4.5: Teachers reported which activities in the AI for Wellbeing curriculum they thought were most engaging and effective for students

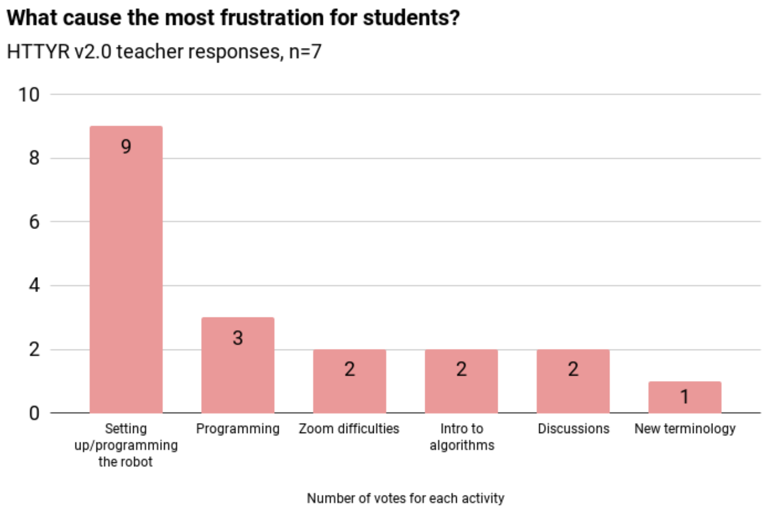

Figure 4.6: Teachers reported which activities in the HTTYR v2.0 curriculum students struggled the most with

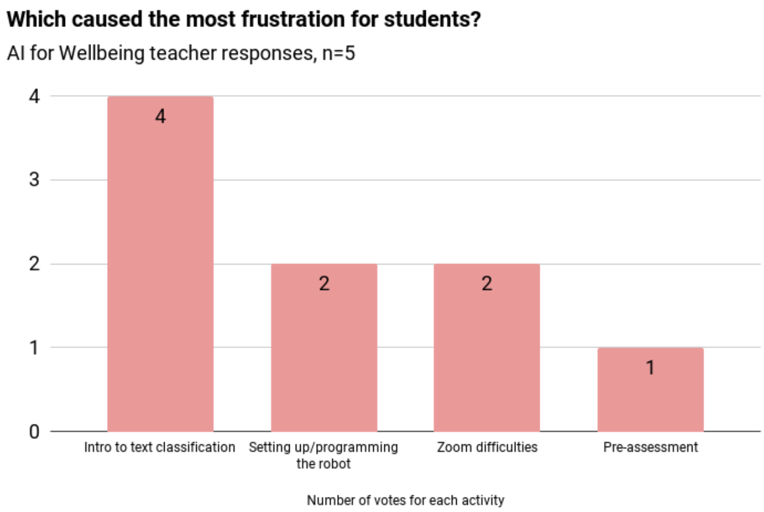

Figure 4.7: Teachers reported which activities in the AI for Wellbeing curriculum they thought were most frustrating for students

may take to support their students effectively through difficult learning moments.

Interestingly, although AI for Wellbeing teachers all indicated that working with the Jibo robot was a powerful point of engagement for students, none of them specifically mentioned the robot in their post-assessment responses about the most engaging part of the course. In HTTYR v2.0, the robots, which were considerably less expensive than Jibo, came up as teachers said they would be interested in using them in future courses if they could come up with the money.

> "I would have to come up with $800 to cover a class. If I want something, I would go buy it (on my own dime)." (P03)
> "We would have more students per robot. . . This is a great opportunity to be like 'you write the code.' 'you do the physical part. Make the robot do what it needs to do.' '[You] talk to stakeholders.'" (P05)

Teachers' Impressions of the Curriculum

Given the overall success of the workshops, teachers in HTTYR v2.0 and AI for Wellbeing were excited to bring the course to their schools in some form. Vazhayil et al. (2019) similarly noted that after teacher training, teachers feel excited to begin using AI curricula in their classrooms immediately [217].

> "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)
> "I absolutely feel more confident in teaching this to the students." (P11)

However, they also acknowledged some real constraints with being able to do that. Across every iteration of the curriculum, teachers expressed that one week was not enough time, especially for new programmers, to practice foundational programming skills and digest new AI information. None of the other schools had an initiative like STEM Week, where they could spend a week working on special projects like AI. It would be up to the teachers to figure out how to rearrange their lesson plans to create space for AI.

The workshop pace was also fast for teachers, who had only a few hours to get used to the new technologies before working on them with students. Teachers felt they would need more time working with tools such as the AI Playground to feel comfortable with them. Finally, there was the matter of money:

> "Schools [in rural areas] are closing and losing funding. . . There are kids out there so how do we serve them?" (P05)
> "I would leave behind the part that costs money." (P02)

As we look forward to broadening access and participation in AI, it is important to reflect on how barriers such as time, training, and lack of resources constrain what teachers can make happen in their classrooms. Teachers with passion and fewer barriers will always be able to implement faster than other teachers. Primarily working with classrooms where implementation is more difficult is essential to ensuring that curricula are designed to reduce teacher implementation barriers.

4.5 Discussion

This chapter described a movement from the theory of the Impact.AI framework to practical curricula that we evaluated with real teachers. Building on what we learned from teachers, we have three key takeaways that can inform the future design of K-12 AI curricula.

Leveraging Three Design Principles for K-12 AI Curricula

In Williams et al. (2022) we introduced the three design principles of active learning, embedded ethics, and lowing barriers to access [112]. That publication primarily discusses how those three design principles led to curriculum designs where students demonstrated technical and ethical understanding, plus the ability to apply their knowledge to projects. With this chapter, we add that teachers' feedback on the curricula also supports the use of these design principles.

[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'

The activities that teachers found most effective were related to the active learning activities that we created, especially engaging in ethical analysis and creating projects. Teachers also felt that embedding ethics was important for students to understand and knowledge that they could take to their other computing lessons. As for barriers to access, teachers both said that AI concepts were well broken down, and that the technology, though not perfect, was something they could bring to the classroom.

The designs of the curricula presented in this chapter and in our previously published works can be effective models for the designs of future AI curricula for K-12 learners.

Training and co-designing with teachers

A unique aspect of our curriculum development work is the level of involvement of K-12 educators. Researcher-educator partnerships are critical to the success of AI, and more broadly, CS education. Goode and Ryoo (2019) discuss three kinds of knowledge that effective computing educators need: technical, content, and pedagogical knowledge [224]. Through researcher-educator partnerships, we can build all three kinds of knowledge. Researchers and teachers need to learn from one another to keep up with AI innovation, effectively translate knowledge to students and navigate the constraints of classrooms.

We believe our researcher-educator partnerships were successful because teachers were interested in using our curricula in their classrooms. Beyond the scope of our research study, we experienced additional barriers when we pursued those implementations, particularly around the structure of a class day and getting teachers the technological resources they needed. It's clear from teacher's feedback and those experiences that accessibility is a concern. We need to continue to iterate on our hardware and software tools to create more options for powerful, yet resource-conscious activities that can work for different classrooms.

AI curricula in different forms

Beyond the one-week model of implementing AI curricula, we also developed AI curricula into the Day of AI outreach program and the RAICA curriculum. Day of AI is another collaboration with I2 Learning. Building off the success of Hour of Code, the goal of Day of AI was to create educational activities that lasted for a few hours that teachers could bring into their classroom as a one-off engagement with AI. Since these activities were a lighter lift and completely online, we have seen thousands of classrooms worldwide engage with them.

The RAICA curriculum was a collaboration between MIT RAISE and Dubai Heights Academy, where we developed semester-long curricula for middle school students. In addition to working in a school in Dubai, the RAICA team co-designed AI lessons with teachers and educators in the United States. Expanding the curricula to a semester model led to more changes in the form of HTTYR than the Day of AI activities. Teachers needed more in-depth lesson plans and ongoing training to prepare for these lessons.

These other forms of AI curricula highlight the importance of flexibility to reach more teachers. It also raised questions about what is important to keep and what can be left behind if the modality of a curriculum changes.

5 Evaluating HTTYR and Dr. R.O. Bott

This chapter evaluates How to Train Your Robot (HTTYR) v2.0 and the AI For Wellbeing curriculum by students' learning outcomes. These data provide evidence that students' conceptual and practical knowledge about AI increased, and that their perspectives about AI further align with a technosocial change agent identity after participating in our curricula.

This chapter discusses:

- The instruments and assessments we used to evaluate students' knowledge and perspectives
- The data we collected with those evaluations
- The implications of these results for the design of K-12 AI curricula

The results of these curricula were previously published in Williams et al. (2022) [112] and Williams et al. (2024) [113]. This chapter includes additional analysis, especially around students' perspectives about AI, and a comparison of these two curricula.

5.1 Instruments and Assessments for Evaluating Student's AI Knowledge and Perspectives

[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'

[113]: Williams et al. (2023), 'Dr. R.O. Bott Will See You Now: Exploring AI for Wellbeing with Middle School Students'

As discussed in the Impact.AI Chapter (Chapter 3, Section 3.1), the Impact.AI learning outcomes can be measured with various assessments that attend to students' developing conceptual knowledge, AI skills, and perspectives about AI. However, there are major limitations in studies that assess K-12 students' knowledge of AI. The assessment space in K-12 AI research is sparse and disconnected as few researchers have developed and shared reliable, validated assessments. Thus, it is difficult to reliably gather evidence about what students have learned or to compare one curriculum or approach to another.

HTTYR v2.0 and AI for Wellbeing use assessments that we continually iterated on, yet these assessments are unreliable. Rather than comparing these two curricula, we use our assessments to gather evidence about what two cohorts of students learned about AI as they engaged with the curricula as a baseline for what is possible for students to learn.

5.1.1 Pre and Post Assessments

To assess participants' familiarity with and perceptions of Artificial Intelligence, students completed a pre-and-post questionnaire on the first and last day of the study. The questions asked about their present familiarity with AI and technology (11 items) and their attitudes toward AI (7 items). The full list of questions can be found in the Appendix.

In AI for Wellbeing, we used a revised version of the original questionnaire. It included more items to assess students' AI knowledge (4 items), a self-assessment of AI practices (1 item with 4 parts), and a survey on students' perceptions of AI (5 items with multiple parts). Some of the items in these instruments were similar to those used in HTTYR v2.0, particularly a question about which everyday technologies used AI. For the remainder of the questions, we sampled questions from previously published AI conceptual examples [141, 163, 225, 226] and perception of AI surveys [211, 226]. The full text of the assessments is available in the Appendix.

Due to the small sample size of student participants and unmatched data, we used the Mann-Whitney U Test to determine whether any quantitative results were statistically significant.

5.1.2 Activity-Based Assessments

We heavily leveraged activity-based assessments in the workshops to gather rich information about what students learned. These assessments were tailored to the content and activities that students did. We analyzed them using mixed quantitative and qualitative approaches, specifically thematic coding to summarize students' responses and rubrics to evaluate the accuracy of students' work. More details about these assessments are presented in this chapter's results (Section 5.3).

5.1.3 Final Project Rubric

Both workshops culminated in a final project where students applied what they learned to passion projects. In HTTYR v2.0, we developed a final project rubric to analyze common themes across students' projects and assess their application of technical and ethical skills. With slight modifications, we leveraged the same rubric to evaluate students' projects in AI for Wellbeing. Since HTTYR focused on supervised machine learning, the rubric contained items about training and test data. AI for Wellbeing focused on design justice and therefore we added rubric items around how well students identified their work's potential benefits and harms.

5.2 Methodology

5.2.1 Online Study Context

In 2020, due to the health emergency created by the COVID-19 pandemic, HTTYR v2.0 had to be conducted online, rather than in the in-person format it was originally designed for. Both HTTYR and AI for Wellbeing were deployed as synchronous online summer workshops, where students used Zoom video conferencing on their personal devices to participate. Students usually had laptops or computers available for the course, but a few had to use mobile tablets. Workshops occurred over the course of a single week (Monday to Friday) with daily sessions that spanned 2-2.5 hours, significantly less than the 30 hours of class time we intended for these lessons.

Other challenges faced in the virtual workshops were maintaining engagement in the class and supporting students. Maintaining the fidelity of the activities in a virtual setting was challenging, given the emphasis on active and collaborative learning in the curriculum design. Some students struggled with their technical setup, lacking resources like working microphones, cameras, or stable Internet connections. Some students were uncomfortable speaking out loud in a new environment. In HTTYR v2.0, all students were just meeting for the first time on that first Monday. In AI for Wellbeing, students had been with one another for another weeklong virtual workshop, but the teachers were all new.

Using robots in the workshops was a significantly ambitious endeavor because it required students to set up their workstations without a trained researcher present to debug. Similarly, when students were working on programming challenges, teachers could not see their students' screens and had to rely on verbal communication to keep everyone successfully moving toward their goals.

We adapted the activities and our teaching styles to address these challenges to accommodate students. Synchronous and interactive tools like Google Classrooms, Google Slides, Padlets, and Kahoots made it easier to facilitate collaboration and engagement. We used breakout rooms and chat-based interactions to have students collaborate, share their ideas, and ask for help. To aid in code debugging, we provided students with illustrated guides and videos they could access outside of class for further support.

5.2.2 Participant Demographics

The study protocol for each respective workshop was reviewed and approved by the Institutional Review Board at the Massachusetts Institute of Technology before we began recruitment. Informed consent procedures involved parents and student participants signing consent forms for the student to participate in the study. Potential participants were assured that they could participate in the workshop without consenting to data collection and they had the right to withdraw their data at any point.

For HTTYR v2.0, we had teachers recruit six middle school students, where we enforced a minimum 50/50 gender balance. In the end, we recruited 26 female students, eight male students, and three students aged 10 to 16 years old (average 12.4) who preferred not to specify a gender. Out of 42 workshop participants, 37 agreed to participate in the study. We did not collect information about students' race or gender to avoid stigmatizing students by collecting such sensitive data. Instead, to ensure racial diversity, we recruited teachers who taught at schools primarily composed of students from marginalized backgrounds. We communicated to teachers our goal of prioritizing the needs of these students and asked that they recruit students representative of their school populations.

For AI for Wellbeing, the program we partnered with collected and shared demographic information about their students. Students ranged

Workshop	Number of students	Age	Gender
How to Train Your Robot v2.0	37	10 to 16 (avg. 12.4)	26 female 8 male 3 prefer not to say
AI for Wellbeing	23	10 to 15 (avg. 11.8)	13 female 10 male 0 other

Table 5.1: Demographic information about student participants from both workshops

in age from 10 to 15 (average 11.8 years old), 13 identified as female, and ten identified as male. Thirteen students (56.5%) identified as Asian American / Pacific Islander, six (26.1%) identified as Black / African American, one identified as multiracial, one identified as Caucasian / White, and two did not specify their race. Thirteen students (56.5%) stated that they spoke other languages besides English at home.

5.3 Results

In this section, we use data gathered during the workshops to answer our research questions about the amount of conceptual knowledge gained and practice with AI skills acquired by students after engaging with our curricula. We will explore this by analyzing students' responses to pre and post-test assessments, activity-based assessments, and their final projects.

5.3.1 Pre-Post Knowledge Assessments

In HTTYR v2.0, the pre-assessment questions were designed to understand the extent to which students recognized how much and when they encountered AI in everyday life. We asked if the following technologies used AI: Google Search, Wireless Printers, Video Calls (e.g., Facetime, Zoom), Game Consoles (e.g., Xbox, Switch), YouTube ("Up next" recommendations), Movie Recommendations (e.g., Netflix, Amazon Prime, Hulu), Snapchat Filters, GPS Apps (e.g., Google Maps, Waze) and Voice Assistants (e.g., Alexa, Sirir, Google Home). In their pretest responses, it was clear that students did not yet have a systematic way of determining if a technology used AI or not. Except for Video Call technology, more than half of students responded that every example technology used AI.

Since students erred on the side of calling items AI, they correctly identified all AI examples, including Google Search, YouTube's "Up next" recommendations, Movie recommendations, Snapchat filters, GPS apps, and Voice assistants. Students were most confident that Voice assistants (100% said "Does use AI"), GPS apps (81%), and Google Search (76%) all used AI. They were least sure about Video calls, Wireless Printers, Game consoles, and Movie recommendations.

On the first day of HTTYR v2.0, students learn to define AI and practice reasoning through whether a technology utilizes AI or not. When asked the question again on the post assessment, students seemed to have a more thoughtful method for determining whether something is AI or not.

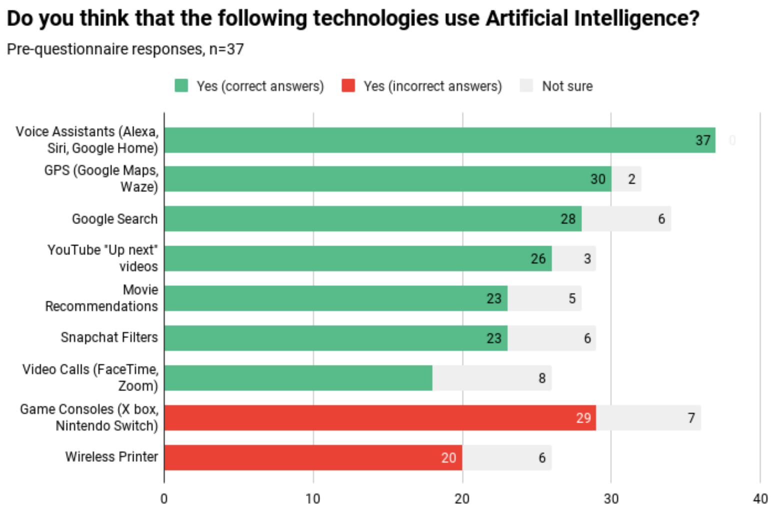

Figure 5.1: Students' pre-test responses to whether several examples of technology used AI or not

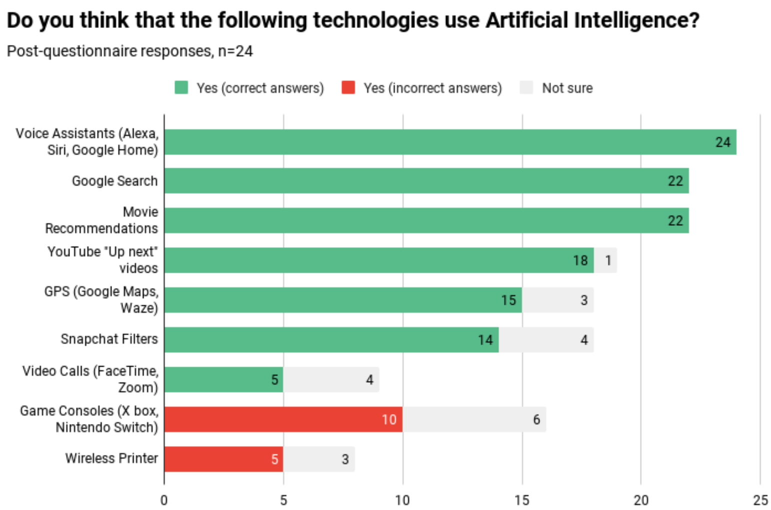

Figure 5.2: Students' post-test responses to whether several examples of technology used AI or not

Some students still used the strategy of claiming that everything uses AI. However, the majority of students correctly determined whether technologies use AI or not. The most confused examples were Video calls (21% said "Does use AI") and Game consoles (42% said "Does use AI"). These results point to the importance of helping students learn to use reasoning to determine whether something is AI, especially as the affordances of AI technologies continually change.

In AI for Wellbeing, we used pre-post assessments with many more questions to assess students' knowledge (Figure 5.3). On the pre-assessment, 19 out of the 23 student participants completed the questions; however on the post assessment, only 10 out of 23 students completed the assessment. On the pre-assessment, we asked four items that assessed AI knowledge, specifically correct definitions of AI, correct and incorrect statements about machine learning, a two-part AI ethics question, and a question about how chatbot systems function. Students' average score on the post-assessment was higher (71%), but not significantly higher than the pre-assessment (61%; U=62.5, p=0.29).

The weakness of the assessment in discriminating high-understanding from low-understanding is evident in the very high pre-test scores, and the reversal effect on the AI ethics question.

The AI ethics question was a two part question where we asked students to identify realistic benefits and harms of AI-powered chatbots, they had to select two of the following answers.

- A. Harm: Chatbots will create too many jobs and there won't be enough people to fill them.
- B. Benefit: Chatbots can talk with people, offering digital assistance whenever they need it.
- C. Benefit: Chatbots rarely make mistakes. They can work perfectly without humans checking them.
- D. Harm: A vehicle that is smart enough to have realistic conversations is likely smart enough to plan a revolt against humans.
- E. Harm: People might put too much trust in robots and use them in inappropriate ways.

Surprisingly on the pre-test most students correctly answered B and E. Given our assumption that students were entering this workshop without prior AI knowledge, this suggests that the answer to the question may have been too obvious. However, on the post-test many fewer students got the answer right. While these data may suggest that students' AI ethics knowledge decreased, their performance identifying potential harms on their final projects combats this interpretation.

A more likely possibility is that the assessment questions was not appropriate for measuring changes in students' AI ethics knowledge. Looking more closely at the data, every respondent selected the correct benefit, that chatbots can be available at any time, on the post-test. However, they chose different potential harms, including that chatbots could create too many jobs and that a self-driving car could become intelligent enough to revolt against humans. Without additional information, it is difficult to understand why students made these selections. But one possibility is that, after discussing different potential AI harms in class students were thinking more expansively about the potential harms of AI and that encouraged them to make these selections.

The largest score increase from pre to post assessment was on the items related to machine learning (pre-test=32% of participants were correct, post-test=64%) and chatbots (pre-test=27% of students were correct, post-test=60%). Therefore, it is possible that these assessment questions were useful.

On the pre and post assessments, we also asked students to self-assess their skills. We saw that students' programming knowledge did not change much from pre-assessment to post-assessment. However, this is not surprising since only one student said they did not have prior programming experience.

There were significant increases in students' self-assessments of their ability to use AI for social good (pre-test median=2, post-test median=3,

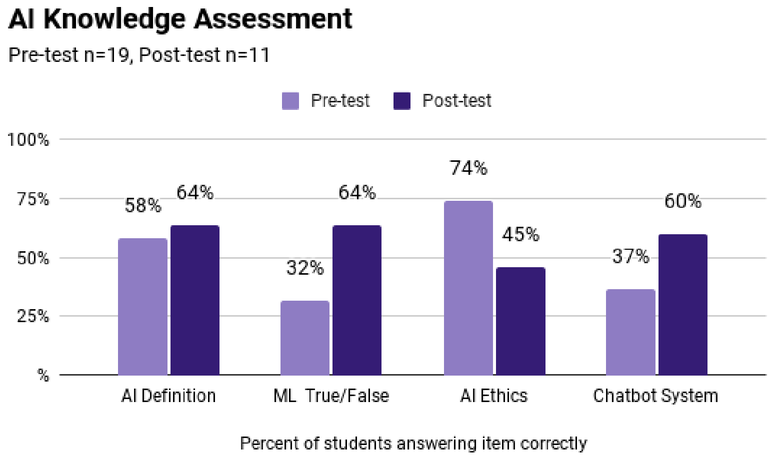

Figure 5.3: Students' pre and post-test scores on the AI Knowledge assessment used in the AI for Wellbeing curriculum

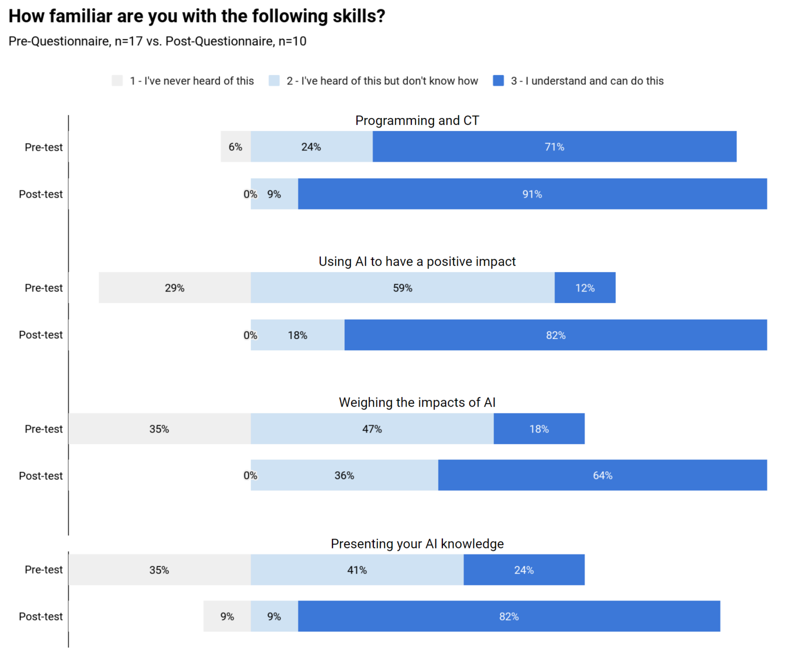

Figure 5.4: Students' pre and post-test scores on the AI Skills self-assessment used in the AI for Wellbeing curriculum

U=23, p=0.015), analyze the ethical implications of AI artifacts (pre-test median=2, post-test median=3, U=38.5, p=0.01), and present about AI (pre-test media=2, post-test media=3, U=39.5, p=0.01). Students gained direct experience with all of these skills during the workshop.

5.3.2 Pre-Post Perspective of AI Surveys

To measure students' perspectives of AI in the first version of HTTYR, we asked them six questions about different facets of their perspectives. Three questions concerned students' affinity for technology, how much they enjoyed learning about it, how much they knew about it, and how interesting they found it. We saw that students rated their affinity for technology highly on the pre-survey and that their perceptions remained high on the post-survey. Similarly, for whether students knew people in the tech workforce, students mostly answered 'Somewhat' or 'Quite a lot.'

However, having a high affinity for technology and knowing people in the workforce did not translate to students wanting to work in computers and technology themselves. Anecdotally, students expressed interdisciplinary interests in medicine, art, and other fields that leverage technology but are not in the computing field. Given that many initiatives for broadening participation in AI are motivated by diversifying the AI workforce, these data suggest that the relationship between students' interest and exposure to technology and their desire to enter the field may be more complex than recognized.

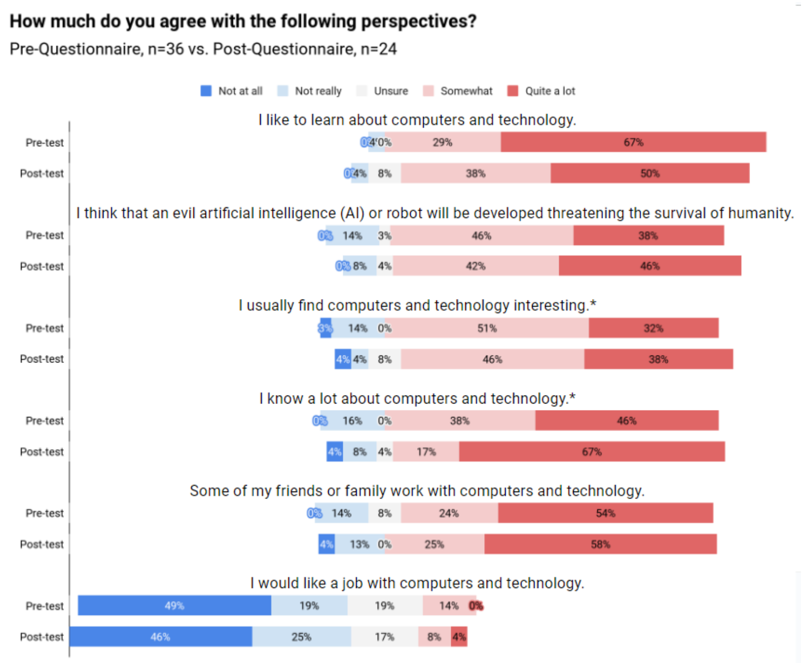

Figure 5.5: Students' pre and post-test scores on the Attitudes toward AI survey used in the HTTYR curriculum

Unrelated to the other survey items, we also asked students how they felt about AI singularity, a superintelligent AI that threatens humanity. Interestingly, many students believed this was possible, and even more so after participating in the curricula. However, no additional questions or anecdotal data further explored this possibility with students. We note it here as an item for future exploration.

In AI for Wellbeing, we used the AI Attitudes survey [226] and MAILS AI literacy scale [211] to create a scale to measure students' motivation to use AI for social good, critical consciousness of AI, self-efficacy, and sense of belonging in the field. Students' pre-assessment scores were very high, and their post-assessment scores remained high. It is important to note that these students and their parents self-selected to participate in this topic, so these data make sense for our sample.

To further explore the data, we conducted an intersectional analysis of the students' survey responses based on where their identities were amongst those that we were specifically working to include and encode in tech. Students who were women, Black / African American, or biracial in our class fell into this category. The differences between Included and Encoded students and the other students were not statistically significant, likely due to small sample size. However, there was an observable pattern in our sample where Included and Encoded students' scores tended to fall on the lower end of the scale compared with their peers.

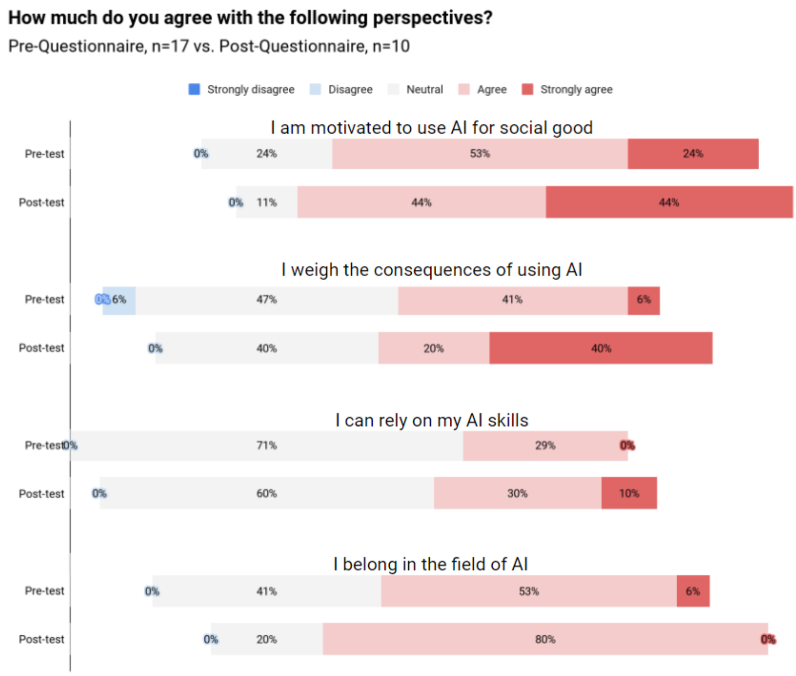

Figure 5.6: Students' pre and post-test scores on the AI Perspectives survey used in the AI for Wellbeing curriculum

None of the Included and Encoded students agreed that they could rely on their AI skills in the pre-survey. Only half of the Included and Encoded students agreed that they belonged in the field of AI, compared to all but one of the other group of students. On the post-survey, we saw Included and Encoded students' data look a lot more similar to the other group. One student scored high on the "I can rely on my AI skills" scale. and all but two fell into the positive end of the "I belong in the field of AI" scale.

Given the small sample size of our data and the fact that our study was a design evaluation and not an experiment, we cannot say that our curricula directly contributed to students' changing perspectives. Still, we believe these data highlight the importance of using intersectional analysis when evaluating AI curricula and not taking for granted that even highly knowledgeable or experienced students from excluded backgrounds will automatically possess a sense of belonging to STEM fields as their counterparts.

5.3.3 Students' Understanding of HTTYR v2.0 Lessons

The main goal of HTTYR v2.0 was to help students better understand image classification, text classification, and the disparate impacts of AI on various stakeholders. Through activity-based assessments, we evaluated students' grasp of the following learning outcomes:

- Students can articulate how neural networks and K Nearest Neighbors algorithms use datasets to learn to make predictions
- Students identify the stakeholders and values relevant to the design of AI systems
- Students create machine learning model projects to address a problem and maximize benefits to stakeholders

Technical Understanding: Identifying Features for Image Recognition

In the neural network activity that students completed on Day 2, they learned about how neural networks learn to classify images by recognizing patterns in the features of training images. As shown in the slide in Figure 5.7, instructors explained how a neural network trained on images of cats and dogs would break down the image into pixels and then use linear combinations of those pixels to recognize features like whiskers, head shapes, and mouth shapes.

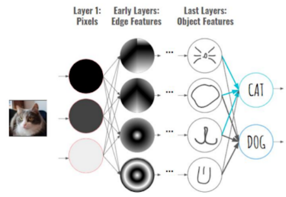

Figure 5.7: Screenshot of a slide used to introduce neural networks for image classification. The slide shows how the image of a cat is broken down into pixels, and then passed through layers of the neural network to recognize high-level features such as whiskers.

Students then practiced training their image classifiers and using them in programming projects. After this hands-on exploration, we asked an assessment question, "What features might a neural network look at to distinguish the suits of cards?" We expected that students who correctly understood how image classification worked would be able to list several distinguishing visual characteristics that can be used to tell apart the shapes of different card suits.

Two researchers graded students' responses, reading through every student's work and marking whether they correctly identified a distinguishing visual feature. Seven students completed this assignment and submitted their work, all of their answers are displayed in Table 5.2.

On average, students listed 4.65 features, 2.82 of which were correct. Students often incorrectly identified features like size, or listed ideas that were too vague (e.g. "one is shaped like a heart") which a neural network would not be able to pick up on.

Ethical Understanding: Determining Key Stakeholders in Real-World Systems

On Day 3, students learned about stakeholder analysis and used an ethical matrix to redesign the Amazon Echo. Originally designed for the home, the students' goal was to redesign the Amazon Echo to be useful in classrooms. At the end of this day, researchers asked students an assessment question to evaluate their ability to identify stakeholders and their values.

> "Amazon is coming out with a smart toaster. A customer will be able to tell their toaster what kind of food they are toasting (slice of bread, bagel, waffle, pizza bites) and it will automatically set the timer and toast their food to their liking. Before selling the toaster to hungry customers all over the world, who are some stakeholders Amazon should consult and what are some issues they might care about?"

Students answered by completing an ethical matrix (Figure 5.8) with three spaces for stakeholders and three spaces for values. We expected students to consider a range of stakeholders, including potential customers, the producing company, and even regulatory bodies for consumer products. We wanted to see values including safety and profit, as these considerations are relevant to developing a new product.

Table 5.2: This table shows the number of students who included each feature of card suits in their responses, n=17

Correct features a neural network might use to distinguish card suits	No. of students
Color Examples: Different color by pattern, colors, red color, black color	16
Sharp points Examples: pointed edges, tips or points, points, sharp edges	14
References to edges Examples: curve shape, inward curves, flat sides	7
Circle shapes Examples: bubbles, circles, bubble placement	3
Other Examples: degrees of angles, indentations, flat bottoms	6

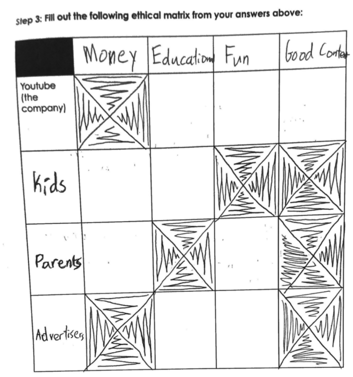

Figure 5.8: Example of an ethical matrix from [100]

We received ten completed ethical matrices, which one researcher graded. The most common stakeholders identified were customers (8 out of 10 responses) and the company developing the toaster (8 out of 10 responses). Students most often identified money or profit (10 out of 10 responses) and the importance of unique features (7 out of 10 responses) as most important for values. These data mirror findings from DiPaola et al. (2020), where students also identified money and popularity as values driving technological design [100].

Considering that this workshop was about responsible AI, students' acknowledgement of the importance of money in the design of AI tools is thought-provoking. During the workshop, students discussed potential conflicts of interest between users' and producers' interests. For example, three out of 10 students identified customers would be more interested in higher quality products while suppliers might prioritize cutting production costs.

One student explicitly pointed out this conflict, noting the tension between what is good for society and sustainable for business. Because we saw students wrestling with this tension both in these questions and in their projects, we decided to offer students a framework for choosing which stakeholders to prioritize in later frameworks. By weighing the relative harms of different design decisions and considering justice, we sought to provide students with tools that helped them recognize and respond to issues of structural inequality. For example in prioritizing users' interests in higher quality products over company's interest in profit especially if doing so meant challenging a harmful cost-cutting practice.

Applying Knowledge: Machine Learning Projects for Social Good

In the last two days, students worked on final projects where they applied their technical and ethical knowledge to passion projects. Students could work on whatever project they wanted using the concepts and tools they learned throughout the course.

Of the 37 student participants, 25 completed and submitted final projects. One pair of siblings collaborated on a final project, so we received 24 projects in total. The remaining 12 student participants encountered significant technical issues with their robots and opted to do a "teach back" as their final instead. Two researchers qualitatively analyzed attributes of the project, including the topic, concepts incorporated, and intended beneficiaries. Then, they evaluated the projects on a rubric that assessed the technical design, ethical design, and implementation.

Fourteen out of 24 projects used machine learning algorithms, eight of the remaining projects used robotics but no machine learning, and one built on a binary decision-making activity we did the second day. Applications of student projects included entertainment (7 projects), helping people (7 projects), healthcare (5), science (3), and education (2). The primary beneficiaries of students' projects were children and teens (5 projects), their families (2), and their communities (2). A gallery of recordings of students' final project presentations is available at `https://httyr.media.mit.edu/afe/afe20`.

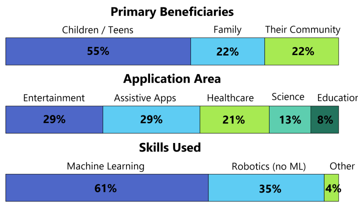

Figure 5.9: The skills used, application areas, and beneficiaries of students' HTTYR v2.0 projects showed the application of machine learning to different genres. Students used machine learning or robotics in entertainment, assistive apps, healthcare, science, and education projects. Students primarily designed projects that benefited themselves, their families, and communities.

Using the following guidelines, two researchers independently rated students' final projects on a scale of 1 to 4 (interrater reliability 0.78). A project met our technical design expectations if it was a well-scoped project that did not propose a task beyond the capabilities of computer algorithms. For students who did machine learning projects, we also wanted them to describe the training data and test data they would use to construct and evaluate their model. A project met our ethical design expectations if students could identify at least three appropriate stakeholders for their project. We asked that students think about how stakeholders might benefit from or be put at risk by their algorithms; we expected them to adequately describe the possible positive and negative impacts of their algorithms. We also expected students to design a user feedback mechanism that could mitigate some of the risks students described. Finally, a project met our implementation expectations if the submitted code ran correctly. If we could debug students' code to make it function, we rated it as "approaching expectations." If the submitted project included a machine learning model, we expected that the model would function correctly and use an appropriate number of training examples.

Students performed well in selecting problems, programming, and constructing machine learning models. An example of an issue with problem selection is a project that tried to use the K Nearest Neighbors algorithm to distinguish the symptoms of a cold, the flu, and COVID-19. The symptoms had too many overlaps for the algorithm to work well.

In model construction, only two text classification projects had issues with their model construction: one had an unbalanced dataset with too many training examples in one class, and the other had fewer than five training examples in their classes. Overall, students' ability to successfully design and build projects with machine learning demonstrates their AI capabilities.

Students struggled the most with identifying a plan to test and improve their applications, sometimes neglecting to design an appropriate mechanism for evaluation and feedback. For example, in a math robot tutor project, the student design mentioned surveying users on how much they

Table 5.3: Distribution of students' scores on final projects in HTTYR

Rubric item	4 Project exceeds expectations	3 Project meets expectations	2 Project approaches expectations	1 Project does not meet expectations
Technical design				
Problem selection, 23 responses total	0	21 (19%)	2 (9%)	0
Identification of training data, 14 responses total	0	11 (79%)	0	3 (21%)
Identification of testing data, 11 responses total	0	8 (73%)	0	3 (27%)
Ethical design				
Identification of stakeholders, 20 responses	0	20 (100%)	0	0
Identification of potential risks, 16 responses	0	12 (75%)	1 (6%)	3 (19%)
Design of appropriate feedback loop, 16 responses	0	11 (68%)	3 (19%)	2 (13%)
Implementation				
Programming, 13 submitted programs	2 (15%)	7 (54%)	4 (31%)	0
Model, 5 submitted text classification models	2 (40%)	1 (20%)	2 (40%)	0

liked the robot tutor but did not include an evaluation based on whether users' mathematics scores improved with tutoring. We discussed the importance of feedback loops in class. Applying these ideas to projects might require more time for reflection, perhaps by engaging with peers or directly with stakeholders during the design process.

This evaluation identified common applications and oversights in students' projects. This information will help us better support students as we use this curriculum in more classrooms.

5.3.4 Students' Understanding of AI for Wellbeing Lessons

The main goal of AI for Wellbeing was to help students understand how, using design justice, social agents like chatbots could be designed to support wellbeing. Individual evaluation of students' work was complex since students worked remotely and time was limited. It was not possible to use planned activity-based assessments, so here we share students' inputs in the class discussions:

- Students can identify the possible harms of AI systems
- Students can use design justice to create social agents for well-being

Ethical Understanding: Recognizing the Benefits and Harms of AI Assistants

On the first day, students completed an Introduction to AI activity that included learning about the benefits and harms of everyday examples of AI. As shown in the slides in Figure 5.10 and Figure 5.11, students learned about six categories of benefits and six categories of harms of AI. The benefits include using AI for healthcare, lifelong education, business, science, social justice and policy, and national defense and security. The potential harms of AI include hurting humanity, manipulating people, inequality and job loss, algorithmic error, algorithmic bias, and violating privacy or compromising security.

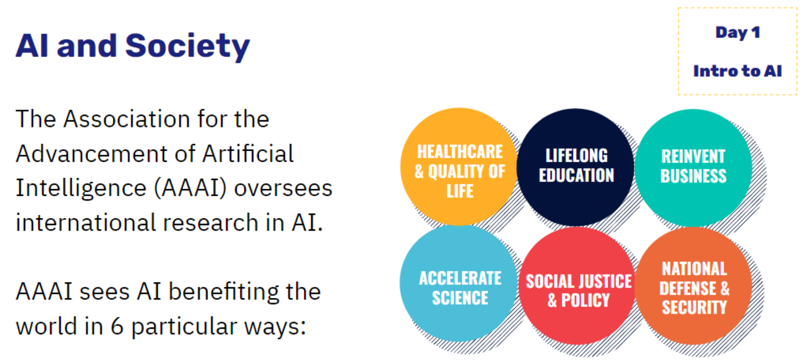

Figure 5.10: Screenshot of slide used to discuss AI benefits

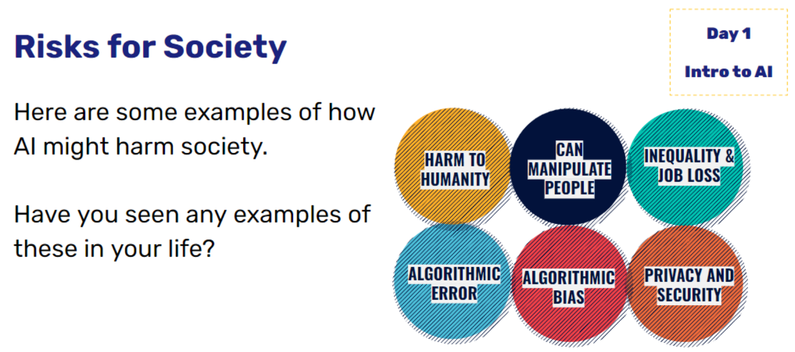

Figure 5.11: Screenshot of slide used to discuss AI harms

After explaining these risks, students thought about whether they had experienced any of these harms in their daily lives. Figure 5.11 shows the responses that students generated through discussion.

Students connected most with examples of AI invading privacy, algorithmic errors and bias, and challenging sustainability. They discussed experiences where AI seemed to violate their privacy or work incorrectly. Teachers prompted students with real-world examples about algorithmic bias and issues of sustainability with AI. Students then connected those broader issues to their own knowledge. The themes that students discussed on the first day became themes that instructors brought up later for students' projects.

Table 5.4: Results of class discussion about everyday AI harms students experienced

Privacy "An example I constantly experience is that whenever me and my friends of family members talk about something specific and there's a phone nearby, later when you go into your phone and open your socials, you see many ads for the same thing you were talking about!"	**Algorithmic Bias** "The test sets that train AI must have diverse sets of data. a lot of AI cannot pick up on facial recognition for darker skin"
Harming Humanity "AI can harm the environment since they need lots of electricity to run those machines."	**Algorithmic Errors and Bias** "ya it cant understand my dads accent"

Applying Knowledge: Designing AI for Wellbeing

Students worked on final projects in the last two days to apply their technical and ethical knowledge to designing their social agents. Students had permission to work on whatever they wanted but were challenged to adhere to the AI for Wellbeing theme and to use the knowledge they learned in the workshop.

Twenty students submitted final projects, which we evaluated with a rubric. Two researchers qualitatively analyzed the content of the projects and then evaluated them on the project's technical design, ethical design, and implementation. Not all students submitted every component of their final projects. If students neglected to submit their code or complete a part of a worksheet, we omitted grading for that item without penalizing them.

One researcher completed the descriptive items evaluation, clustering projects by their application space and primary beneficiaries. Figure 5.12 shows that the 20 student projects covered a wide range of application areas and targeted many different beneficiaries. Six projects applied to the medical field, including many recreations of the personal health robot Baymax from the children's movie Big Hero 6. Six projects related to mental health and included robot companions and an embodied version of the Headspace app. Two projects were robot companions for one, one of which was specially designed for kids by kids. Another notable project application sought to advance the field of conversational AI in general by improving AI's ability to engage emotionally with people. A gallery of students' projects is available at `https://sites.google.com/media.mit.edu/dr-robott/sa23`.

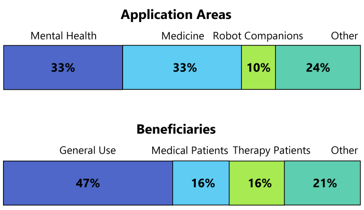

Figure 5.12: The themes and beneficiaries of students' AI for Wellbeing projects showed students applying their knowledge to a diverse range of projects. Students' projects applied to medicine, mental health, and robot companionship. Almost half of the projects were meant to benefit the general public, while others specifically helped medical patients, therapy patients, nurses, kids, non-English speakers, or the project's author.

To evaluate the effectiveness of our project-based teaching method in helping students apply AI concepts to their own lives, we developed a rubric for students' final projects. The rubric evaluates projects 1) technically on their relevance to AI, problem definition, and context definition, 2) ethically on identifying stakeholders and potential positive and negative impacts, and 3) on programming. Two researchers independently used a one-point rubric to evaluate students' work; students either received one point for accomplishing the objective or they received zero points.

Using this one-point rubric increased the interrater reliability to 0.95. We provided students with a rubric to help them keep track of these goals.

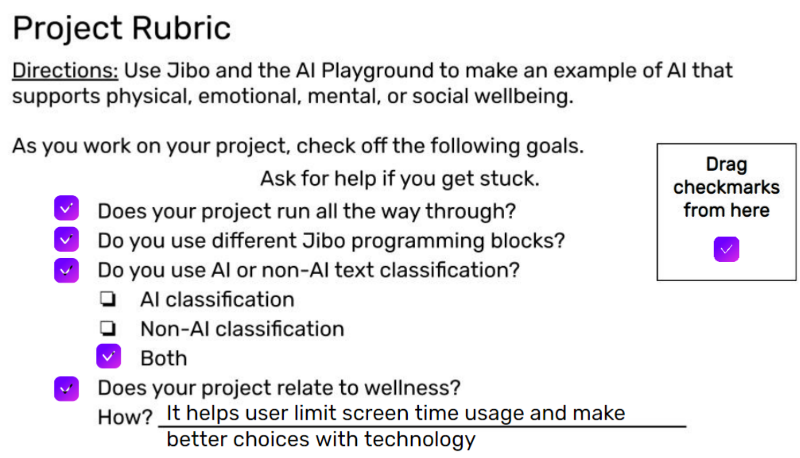

Project Rubric

Directions: Use Jibo and the AI Playground to make an example of AI that supports physical, emotional, mental, or social wellbeing.

As you work on your project, check off the following goals.
Ask for help if you get stuck.

- [x] Does your project run all the way through?
- [x] Do you use different Jibo programming blocks?
- [x] Do you use AI or non-AI text classification?
  - [ ] AI classification
  - [ ] Non-AI classification
  - [x] Both
- [x] Does your project relate to wellness?
  How? It helps user limit screen time usage and make better choices with technology

Drag checkmarks from here

Figure 5.13: The rubric shared with students to guide their work on the final project

A project met our technical design expectations if it was an AI-relevant, well-scoped project that did not propose a task beyond the capabilities of computer algorithms. A project met our ethical design expectations if students identified at least two stakeholders and at least one positive and negative impact for their project. Finally, a project met our implementation expectations if the submitted code ran correctly, used the programming concepts we discussed in class, and was sufficiently adapted from any provided code examples. Table 5.5 shows the distribution of students' scores on the final project.

Students very successfully designed AI-relevant projects that correctly ran on their robots. Nineteen students submitted the ethics portion of the project and were moderately successful at analyzing their ideas. Eleven projects identified at least one relevant stakeholder for their project. Only ten students filled out the positive and negative impact portions of their worksheets but, of those, eight identified at least one potential harm of their work based on concepts learned in class. We would have liked to see more students complete this portion of the worksheet since identifying the negative consequences of your work is an important skill. Eleven students submitted their project code, and all eleven projects ran successfully.

Notably, five projects were based on the exemplar project code we provided students. In the AI for Wellbeing project, we provided example code for projects, acknowledging previous students' struggles completing functional projects. Given the time constraints of the workshops and the amount of new programming skills students were learning, previous students faced significant challenges in building their desired projects from scratch. The ability to remix existing code is a valuable construction skill which, for our purposes, was a practical option to offer our students. Therefore, although several projects closely resembled the exemplar projects, likely due to limited time and students' lack of familiarity with the AI Playground, students did successfully add personal touches to align their projects with their project goals.

Table 5.5: This table shows the number of students who accomplished each rubric item on their final projects in AI for Wellbeing, n=17

Rubric item	No. (%) of students
Technical design	
Relevant to AI, 17 responses total	17 (100%)
Appropriately scoped project description, 17 responses total	17 (100%)
Meaningful project goal, 17 responses total	16 (94%)
Ethical design	
Identified at least one relevant stakeholder, 17 responses	11 (65%)
Identified at least one positive impact, 11 responses	11 (100%)
Identified at least one negative impact, 10 responses	8 (80%)
Implementation	
Runs through correctly, 11 submitted programs	11 (100%)
Uses programming concepts from class, 11 submitted programs	10 (91%)
Unique from provided examples, 11 submitted programs	6 (55%)

5.4 Discussion

This chapter discusses how we evaluated students' grasp of AI concepts, practices, and perspectives before and after completing the How to Train Your Robot and AI for Wellbeing curricula. Noting that there are significant challenges with assessing students' learning in K-12 AI curricula do to the lack of standardized evaluations, we believe our data illustrate promising results for our curricula. Here we center two takeaways for the design of future Impact.AI-aligned curricula and Including and Encoding all students.

Designing Impact.AI-Aligned Curricula To Support Students' Developing Knowledge and Perspectives

[112]: Williams et al. (2022), 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'

Williams et al. (2022) [112] discussed the importance of our three design principles to students' demonstration of technical and ethical AI knowledge. The AI for Wellbeing curriculum data do not contradict what we saw in previous work. Students' final projects, where they applied their expertise to personally meaningful projects, demonstrated how well students had mastered the subject content.

The Impact.AI framework offered an additional strength to the design and assessment of our curricula because it standardizes expectations. The Impact.AI framework guided the development of our learning objectives and assessments in AI for Wellbeing, making it easier to justify activity designs and interpret outcomes. A big goal of the AI for Wellbeing curriculum was to help students realize they could use their

AI knowledge for social justice. This goal led to an activity design where students learned about design justice and discussed how the harms of AI were impacting them personally. This discussion on the first day directly influenced some of the projects students presented on the final day, demonstrating their technosocial agency.

In Chapters 3 and 4, we explored how using a common competency framework can make it easier to compare the designs of different curricula. This chapter demonstrated how a common framework was also essential in evaluating curricula. If translated into rubrics and assessments, frameworks like the Five Big Ideas in AI and the UNESCO AI Competency framework are best positioned to offer standardization to K-12 AI instruction work. Though we advocate for a more culturally responsive framework like Impact.AI, creating a standard to unite future K-12 AI instruction efforts would be beneficial.

Including and Encoding Students With Culturally Responsive-Sustaining Pedagogy

The perspective assessments we used in AI for Wellbeing, particularly the AI attitudes survey, demonstrated that participating in AI curricula can make a difference in students' AI identities, which matches prior research findings that engagement in STEM activities at a young age is generally beneficial for students [14]. However, our results also highlighted the importance of attending to the STEM identities of students who have been historically excluded or excoded from tech. Using frameworks like Impact.AI, curriculum designers might be better equipped to intentionally boost these students' perspectives about themselves and AI, which is essential given the complex barriers these students face. In the future, we hope to see more K-12 AI curricula leverage CRPs and other educational equity frameworks with larger-scale experimental studies that collect evidence around which design approaches make the most difference.

[14]: Sullivan et al. (2016), 'Girls, boys, and bots: Gender differences in young children's performance on robotics and programming tasks'

6 Designing Tools for AI Education

In this chapter, we describe the design of Sparki, a generative AI chatbot that provides feedback and design support on users' AI projects. We designed Sparki for use in Impact.AI-aligned curricula, or any project-based AI curriculum where students apply their technical and ethical knowledge of AI to open-ended projects. As students work on their projects, they can interact with Sparki to deepen their thinking about ethical considerations and get assistance programming in the RAISE AI Playground. Through user studies, we evaluated Sparki's usability and effectiveness in supporting users working on AI projects. This chapter discusses:

- A framework for designing AI education tools that scaffold open-ended projects
- The design of Sparki, an interactive chatbot for K-12 AI instruction
- The evaluation of Sparki

Considerations for ongoing work in designing AI education tools to support K-12 students We believe this tool will help make it easier for AI and programming novices to learn about AI as they construct open-ended projects.

6.1 Motivation for Sparki

AI courses often leverage technology such as programming platforms and intelligent tutoring systems to engage learners with hands-on opportunities to learn about AI through construction and design [44]. These experiences can make AI concepts more tangible and concrete. However, they also create challenges in the classroom. K-12 students are often new to computational thinking and programming. In Chapter 4, teachers often gave feedback that students needed support when designing and implementing their projects, especially in short interventions.

We used teachers' feedback to create a framework for designing AI tools to scaffold learners working on open-ended AI projects. Figure 6.1 shows how AI tools might provide technical, metacognitive, and reflection scaffolding [227]. Technical support includes:

- Engaging students with interactive tutorials
- Offering continuous feedback as users work on projects
- Just-in-time support when users need help
- Facilitating formative assessment

Metacognitive support, especially for promoting creativity, includes:

- Modeling divergent thinking
- Offering positive reinforcement during and after project work time
- Prompting users to engage with their creativity

In this work, we will focus on those two since those are the kinds of scaffolding Sparki provides.

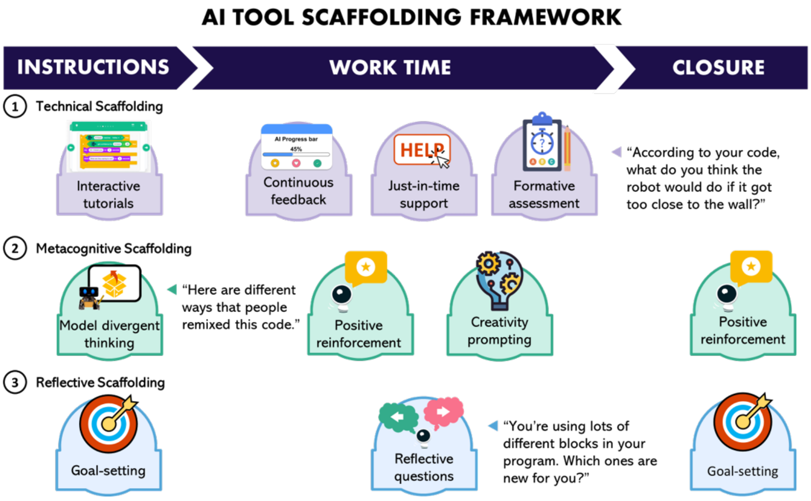

Figure 6.1: The AI Tool Scaffolding Framework delineates the ways that an AI tool might support student working on open-ended projects through technical, metacognitive, and reflective scaffolding.

This framework informed the design of two AI education tools we integrated into our curricula, LevelUp and Sparki. LevelUp was an automatic code assessment tool we built into the programming platform students used to work on their AI projects to support their work by giving feedback [225]. In an evaluative study of the tool, we saw that having access to LevelUp led to significant improvements in learners' projects. Even though users enjoyed the tool, they still requested more personalized guidance and the ability to ask the support tool for more support or details.

More ML education tools should be built into the classroom experience to support students' project creation. In prior work, AI tutors and robot mediators have effectively scaffolded and supported students' learning [225, 228]. Learning agents have been shown to help students develop creativity and self-efficacy [103, 105, 106]. Engaging in reflection with support from an agent can enhance students' creative process and understanding of how conceptual knowledge can be applied [107, 108]. Automatic scaffolding tools can provide technical, metacognitive, and reflective scaffolding for students and educators as they learn [227].

6.2 System Description

We designed Sparki as a creative companion that provides technical and creative scaffolding as students work on open-ended AI projects. Sparki's technical scaffolding features include making tutorials more interactive, offering continuous feedback on students' work, and providing just-in-time support. Sparki leverages GPT3 to evaluate users' input against a rubric, generate code, and generate responses to user's questions. Sparki also models creative thinking and offers positive reinforcement. Sparki's modeling and positive reinforcement behaviors are controlled by a combination of pre-written and GPT3-generated examples.

We also designed Sparki with design and safety considerations front of mind. One of the most pressing concerns for generative text applications is overreliance. We wanted to avoid having users over-rely on the

GPT3 algorithm's responses, which are known to generate inaccurate information. We made Sparki's responses more reliable by anchoring its generated responses to factual information in an AI knowledge database we compiled for this use case. Another consideration for generative text applications is preventing intentional or accidental misuse. To prevent Sparki from being used as a tool to generate harmful, incorrect, or off-topic language, we constrained Sparki's generations to short discourses and instructed it to only discuss AI and programming.

Given that our target audience was minors, we also took steps to protect users' privacy. Before interacting with the system, we warn users not to input private information into Sparki. Although Sparki utilizes the GPT3 API, which claims to not collect personal information, we wanted users to be aware that they were interacting with a 3rd party system. We also limited the use of Sparki to the bounds of class time. During the workshop, we showed students how to use Sparki, explained how text generation worked, and then monitored their use.

6.2.1 Interface Design

There are different forms of Sparki: a chatbot interface, a version built into a brainstorming worksheet (the AI for Wellbeing curriculum used this one), and a social robot interface built into the RAISE AI Playground.

Sparki Chat Interface. The original form of Sparki is its chatbot interface. Leveraging an open-source library called react-chatbot-kit[1], this interface contains a chat window, input box, and modules that users can interact with to communicate with Sparki. The live version of this interface is available at `https://mitmedialab.github.io/sparki`.

1: react-chatbot-kit Library, N.d. [Github repo], `https://github.com/FredrikOseberg/react-chatbot-kit`

The chat window and input box replicate a normal chat window that one might use to communicate with other people. This interface is familiar and easily navigated by most users.

The chat window has a few kinds of built-in modules. The first module is a set of buttons with pre-written questions the user might want to ask Sparki. These buttons are displayed in Figure 6.3, and have different information depending on what the user is working on, such as"Show me an example" "Give me some ideas" and"Progress check." These buttons help scaffold interactions with Sparki, showing them examples of effective questions they might ask to work with Sparki.

Another important module for the chat interface is the Scratch blocks module. GPT can generate code examples based on what it learned from the pieces of the Scratch wiki that were part of its dataset. Using the scratchblocks[2] package, we can render the raw text into images of Scratch blocks, as shown in Figure 6.2. This creates a much more pleasant interface for discussing Scratch code with Sparki.

2: scratchblocks library, N.d. [Github repo], `https://github.com/scratchblocks/scratchblocks`

The final module in the chat interface is a pop-up modal of cards that are also shown in Figure 6.2. The module is designed such that decks of images and text, including images and text generated on the fly, can be displayed in a persistent window that users can scroll through. Currently, we only use this module to display static tutorials and information.

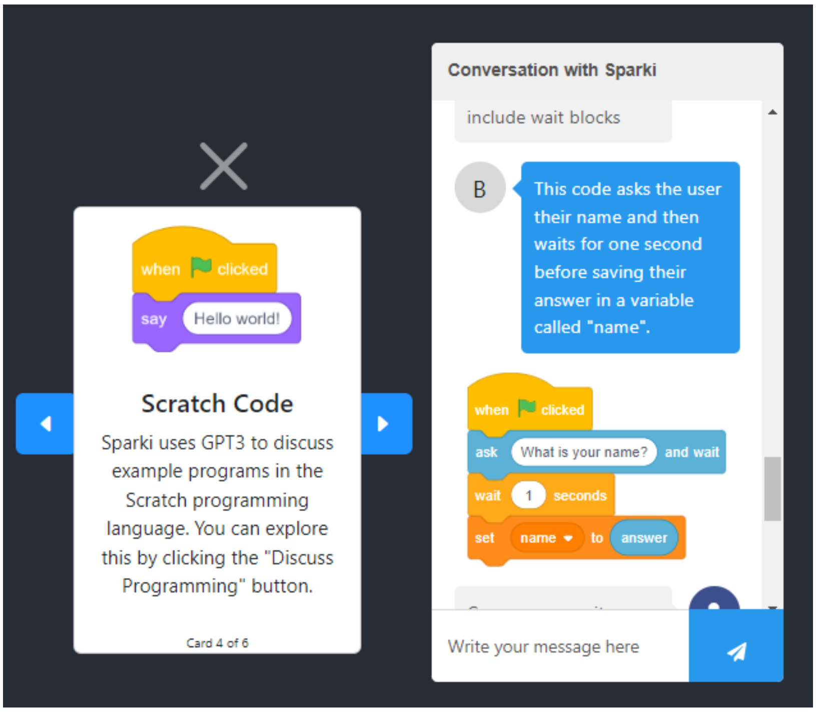

Figure 6.2: Chatbot interface of Sparki showing the Scratch blocks code and card modal modules

Sparki Project Form. The next form of Sparki is the project form interface with input boxes where users write down the project title, description, stakeholders, and positive and negative impacts of their project idea. This version of Sparki is available live online at `https://mitmedialab.github.io/sparki/ai-project`.

The Sparki chat interface is built into every textbox so users can access it as they write their project proposals. Like the original Sparki chat interface, each interface attached to the project input boxes has buttons with pre-written questions. The questions are different for each input box.

The organization of the sections and questions scaffold users through a design justice brainstorming process. Design justice is a process for creation and redesign that centers people who are usually marginalized [28]. We chose to center design justice in the design of the Sparki project form because it is a key learning objective in the AI for Wellbeing curriculum that the form was originally designed for.

Another unique module of the Sparki project form is the progress check module. This module, shown in Figure 6.4, automatically updates to give users feedback on how well their proposal matches the rubric shown in Table 6.2.

Sparki social robot interface. The final form of Sparki is a version built into the RAISE AI Playground and the social robot, Jibo. As shown in Figure 6.5, users can converse with Sparki about their projects through Jibo, which is also connected to the RAISE AI Playground. This version of Sparki has access to users' code and the tutorial cards built into the Playground. This makes it possible for the agent to contextualize its responses to the screen users are looking at. Jibo responds to a user's question using information from the tutorial card shown on the user's screen.

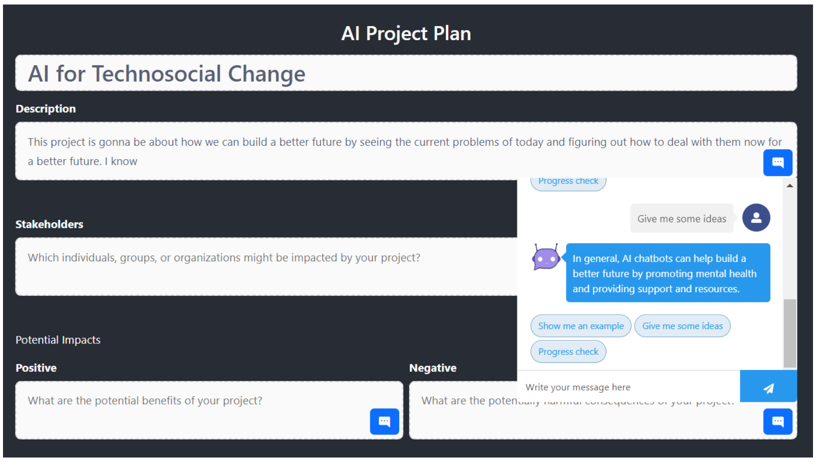

Figure 6.3: The Sparki chatbot is built into each section of the project proposal form. Users can select one of the pre-written questions or type into Sparki's chat input box

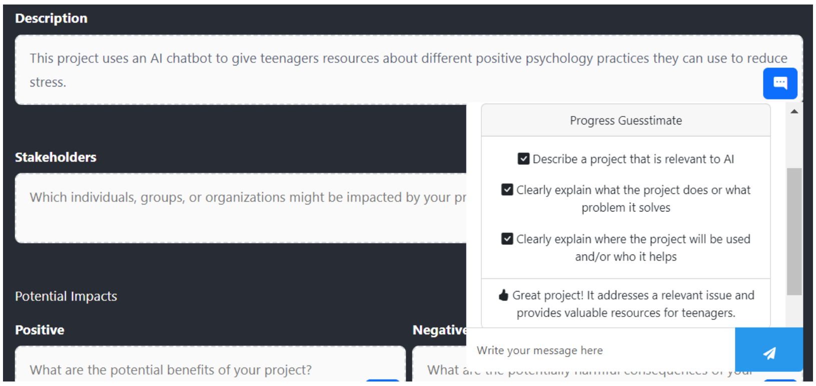

Figure 6.4: The Sparki chatbot contains a progress check module that automatically updates as the user works

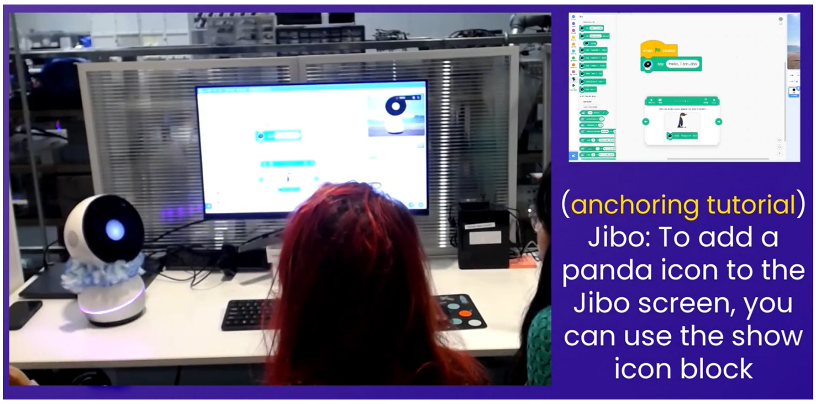

Figure 6.5: In the RAISE Playground, Sparki controls the Jibo robot to answer users' questions as they work on projects

Being built into a social robot makes it easier for more than one user to engage with Sparki at once. As discussed throughout human-robot interaction literature, the use of embodied agents to socially mediate learning offers significant benefits, especially when learners are working in group.

Design considerations. Like other Impact.AI-aligned platforms, Sparki was designed with access in mind. Sparki is a free tool available on the Internet. However, users need Internet access and an OpenAI API key to use the site. Since this part of Sparki costs money, we provide workshop participants with a free API key. In the future, we should design a free version of Sparki that does not rely on Internet connections.

All versions of Sparki work even on low-resourced devices, as long as they have a browser. The chatbot version of Sparki is also compatible with mobile phones. Additionally, Sparki is compatible with screen readers and voice input tools for low-vision users, users for whom reading and writing in English is difficult, and anyone else who would benefit from these alternative interaction mechanisms.

6.2.2 System Architecture

In the remainder of this chapter, we will focus on the Sparki Project Form, as that is the version we used in curricula and evaluated with user students. As shown in Figure 6.3 and Figure 6.4, the project form has three major components: the project form interface, Sparki chat interface, and the AI knowledge base.

The project form interface, which we described in the previous subsection, is one of the main ways that users interact with Sparki. Every time the user changes something on the form, a state variable is updated that sends a report to Sparki. The next main component, and the other main interface for communicating with Sparki, is the chat window. Every time the user types something into the chat window, that input is stored as user input.

The final component of Sparki is the AI Knowledge Base. The knowledge base contains definitions, examples, and other information about each project section. We compiled the knowledge base from several AI and ethics resources, including the AI Audit Game [229] which describes common harms of different applications of AI. We also used the AI Development Canvas [230], a project development worksheet very similar to our project interface, that guides designers through thinking through the implications of their AI projects. Finally, we used the Ethics and Algorithms Toolkit [231], another worksheet tool that guides designers to think through the potential risks of their work.

[229]: Ali et al. (2023), 'AI Audit: A Card Game to Reflect on Everyday AI Systems'

[230]: Dignum et al. (2020), *Tools to operationalize the UNICEF policy guidance on AI for children*

[231]: Anderson et al. (2021), *Ethics and algorithms toolkit*

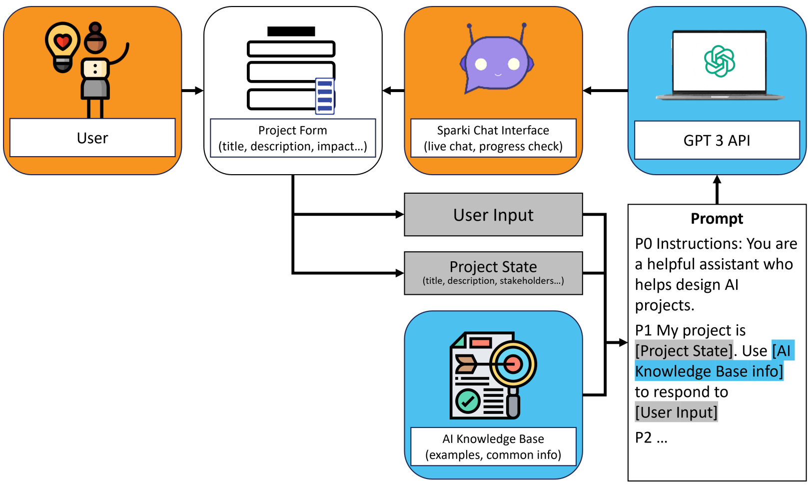

Figure 6.6: A diagram of the Sparki Project Form's system architecture

Whenever the user inputs a new prompt in Sparki's chat window, Sparki returns a response from GPT3's API. Behind the scenes, Sparki combines the user's input, the current project state, and information from the knowledge base to create a prompt to send GPT3. This is done to give GPT3 context for the user's input and to reduce the potential for hallucination. The potential for inaccurate and inappropriate responses is a major concern for all generative AI tools. By anchoring GPT3's response to accurate information from the knowledge base, we attempt to control its generations. We explicitly instruct GPT to respond to any user's questions based on information from its knowledge base. Plus, when Sparki first begins, we provide instructions for GPT3 to respond as a friendly, helpful agent for working on AI projects. We instruct it to keep its responses short and to not allow even hypothetical deviations from conversations about AI.

Section	Project description	Stakeholders	Positive impacts	Negative impacts
Default text	What does your project do?	Which individuals, groups, or organizations might be impacted by your project?	What are the potential benefits of your project?	What are the potentially harmful consequences of your project?
Initial chat message	The project description should describe what problem your project solves and for whom.	Stakeholders are anyone interested in your project and its outcomes. Can you think of at least three people who might be impacted by your project?	Positive impacts are ways that your project helps your stakeholders.	Negative impacts are ways that your projects might (unintentionally!) harm your stakeholders or put them at risk.
Pre-written questions	Show me an example, Give me some ideas, Progress check	Show me an example, Who are vulnerable stakeholders, Progress check	Show me an example, Give me some ideas, What is a high impact project, How do I use design justice, Progress check	Show me an example, Give me some ideas, What is a high risk project, How do I use design justice, Progress check

Table 6.1: Each section of the Sparki form interface has a default input prompt, an initial message from Sparki, and pre-written questions in the chat window

Table 6.2: After users completed their AI project proposals, we graded them on each section of the proposal using this rubric.

Section	Project description	Stakeholders	Positive impacts	Negative impacts
Rubric Items	1 Describe a project that is relevant to AI	1 List at least three, specific stakeholders relevant to this project	1 List at least three, specific positive impacts of your project	1 List at least three, specific negative impacts of your project
	2 Clearly explain what the project does or what problem it solves	2 Name at least one stakeholder that might be more vulnerable than others	2 Mention who is most positively impacted (hint: it should be a vulnerable stakeholder)	2 Mention who is most negatively impacted (hint: try to protect vulnerable stakeholders)
	3 Clearly explain where the project will be used and/or who it helps		3 State whether your project is low, medium, or high impact	3 State whether your project is low, medium, or high risk

6.3 Methodology

We measured the effectiveness and usability of Sparki through evaluative user studies. We conducted two user studies, one pilot with secondary students who used the tools in AI literacy workshops and one experimental study with adults.

6.3.1 Participants

Pilot Study

Through the pilot studies, we collected initial feedback about Sparki from the target audience of the tool. We recruited participants from two separate summer camps where they were exploring AI and design. In accordance with the consent procedures dictated by MIT's Internal Review Board, we first got consent from all student participants and their parents. Students were not required to consent to having their data collected to participate in the program. A total of 22 students participated in the study ranging in age from 10 to 17 years old (average 14.5). Eight (8) students identified as male, seven (7) as female, one (1) as non-binary, and six (6) neglected to say. All students participating in this study had prior knowledge of AI and design justice because they were part of their workshops.

Experimental Study

With the experimental study, we wanted to further understand which of Sparki's affordances a difference in the quality of users' AI proposals. Since this experiment was only a first step, we recruited a small sample of 10 participants (Table 6.3) using the online study recruitment tool Prolific.com. We randomly assigned participants to conditions that dictated the order in which they were exposed to Sparki as they worked on project proposals.

Since participants were randomly assigned, we could not control the size of the condition groups or counterbalance them based on demographic information. The majority of the participants were male (6 out of 10), the other participants included 3 females and one male. About half the participants were below the age of 30, while the other half were above the age of 30.

		Gender			Age Range			
Condition	Total	Fem.	Male	NB / Third	18 to 24	25 to 29	30 to 39	40+
A (w/o Sparki first)	6	1	4	1	2	0	3	1
B (Sparki first)	4	2	2	0	1	2	0	1
All	10	3	6	1	3	2	3	2

Table 6.3: Gender and age range of participants

6.3.2 Procedures

Pilot Study

In the user study with students, everyone had around 40 minutes to use their AI knowledge to brainstorm project ideas on the Sparki project form. After students were done, they completed a feedback form. We used participants' feedback to improve the Sparki interface.

Experimental Study

The experimental study used a cross-over study design where participants completed two project proposals. In Condition A, students first did their projects without Sparki's AI help. As shown in Figure 6.7, they used the same project form interface as the one with Sparki, but instead of having a chat window with access to Sparki, users had a static information box with all of the information from the Knowledge base. Figure 6.8 shows a diagram of the study procedure we used in the experimental study.

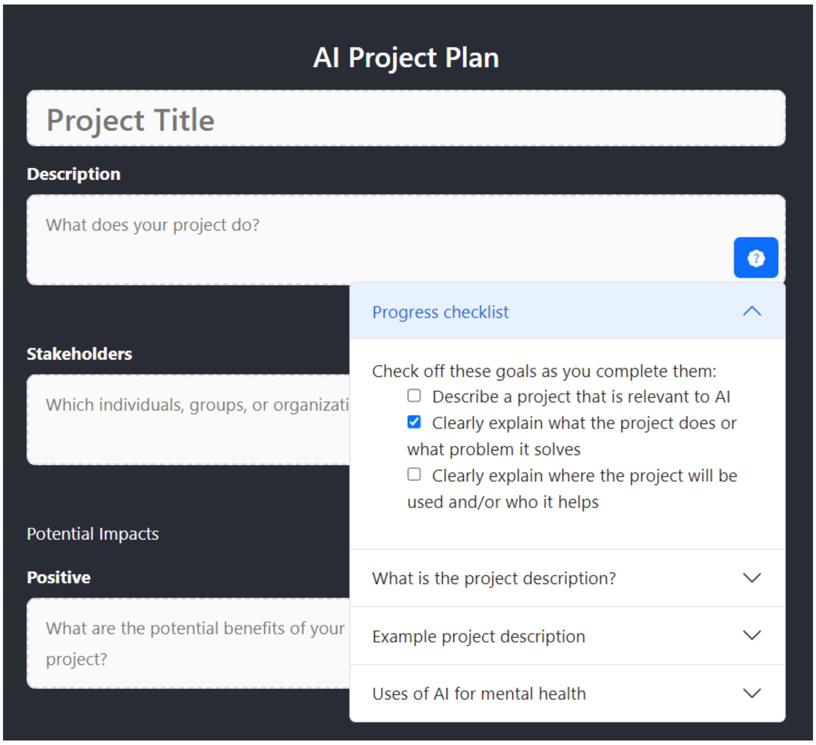

Figure 6.7: The Sparki Notes interface that participants used on their projects when they did not have access to Sparki.

1. Pre-study survey. First, participants created a pre-questionnaire about their prior experience with AI, their preferred method of working (individually or collaboratively), and their expectations of an AI writing tool. At the end of the survey, users watched a video tutorial introducing them to the project writing activity and showing them how to use Sparki.
2. Proposals. Next, participants completed their first project proposal. The theme of the first project proposal was "AI for Wellbeing," and the second was "Self-Driving Cars for Social Equality." The examples and information built into Sparki's knowledge base related to the proposal theme. When participants completed all

five sections of the project form, they submitted their final proposal and logged their interactions to a Firebase server. Participants in Condition A used Sparki notes on their first proposal, while in Condition B, they used the AI version of Sparki first. For the second project, the two conditions switched.

3. Post-activity surveys. After completing each project proposal, participants answered questions about the tool they had just interacted with.
4. Post-study survey. At the end of both proposals, participants completed final questions about the tools they had interacted with and their experience completing the activity.
5. Proposal rubric. We graded participants' projects using a single-point rubric that corresponded with the rubric built into Sparki. Both human and computational graders evaluated students' final submitted proposals. For the human graders, we had them rate the proposals independently, and then compare their responses. We used this information to calculate interrater reliability. For the GPT grading, we gave the following sequence of prompts to ChatGPT-3.5-turbo.

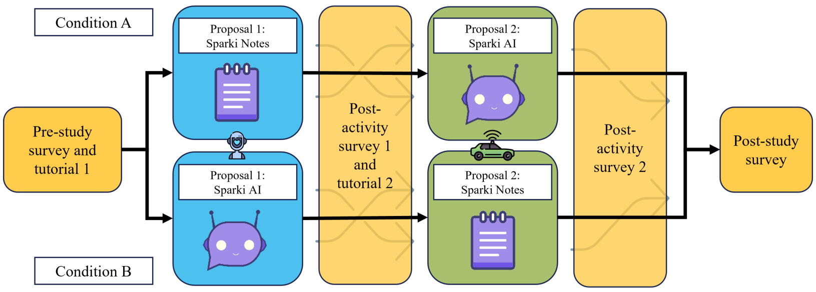

Figure 6.8: In the evaluative study, users were divided into two conditions. Each condition completed one project proposal with Sparki and one without it.

The full text for all of the questionnaires and the rubric is available in the Appendix.

Table 6.4: GPT prompts given to Sparki to grade project proposals

1 Act as an expert educator who can read information, analyze text and give supportive feedback based on a rubric that I will give you. Students were given the following assignment: Create a proposal for a project that uses AI to help with education or maintaining well-being. Describe the project you want to create, the main stakeholders you should consider, and the potential positive and negative impacts of the work. When you are ready, I will give you the rubric. I will then begin to give you the student work to evaluate, and you will provide a numerical grade for the work. Do you understand?
2 Rubric: 1.1 Award one point if the project is relevant to using AI to solve a problem. Put zero points if you don't think the problem has anything to do with AI. . . .1.2 Award one point if what the project does or what problem it addresses is clearly explained. Put zero points if the problem description is very broad ("it will make the world better") or very vague ("it is about animals")....Are you ready to start grading?
3 Give a score to this project proposal. name: [PID]; description: [paste from proposal]; (repeat for all proposals)
4 Put the last 22 scores into a table, and sort alphabetically by name.

6.4 Results

6.4.1 Pilot Study Results

We rated participant's final project rubrics using two human graders and ChatGPT. The interrater reliability between human graders was satisfactorily high, at 0.95. With ChatGPT, the interrater reliability fell to 0.73. In cases where there were disagreements in ratings, we went with the majority score. Any ties were broken by having the human graders discuss until they agreed on a final score.

Since there was nothing to compare scores to in the pilot study, the results of the grading can only serve as a baseline for setting expectations. As shown in Table 6.5, participants were most successful in describing the problems and context in which they wanted to apply their work in the project description. They struggled more with ensuring their ideas were relevant to AI. On average, users identified 1.75 stakeholders, 1.7 positive impacts, and 1.25 negative impacts. We would have liked to see students consider more stakeholders and have a better balance of identifying potential positive and negative impacts.

Table 6.5: Average final project scores in the pilot study

Rubric item	Average score
1.1 Relevant to AI	0.76 (of 1)
1.2 Defines problem	0.81 (of 1)
1.3 Defines context	0.86 (of 1)
2.1 List of stakeholders	1.75 (of 3)
3.1 List of positive impacts	1.7 (of 3)
4.1 List of negative impacts	1.25 (of 3)
Total	7.13 (out of 12)

User feedback on Sparki was overwhelmingly positive. Five (5) out of 7 participants who completed the post-study Agreed or Strongly agreed that Sparki was exciting or engaging and made their task more manageable. Six (6) out of 7 said that people could learn to use Sparki quickly and that they felt confident showing someone else how to use it. Five (5) out of 7 students disagreed or strongly disagreed that Sparki was difficult to use.

As shown in Figure 6.9, users reported using Sparki in a variety of ways, including coming up with new ideas (6 out of 7), exploring alternative ideas (4 out of 7), and overcoming creative blocks (4 out of 7).

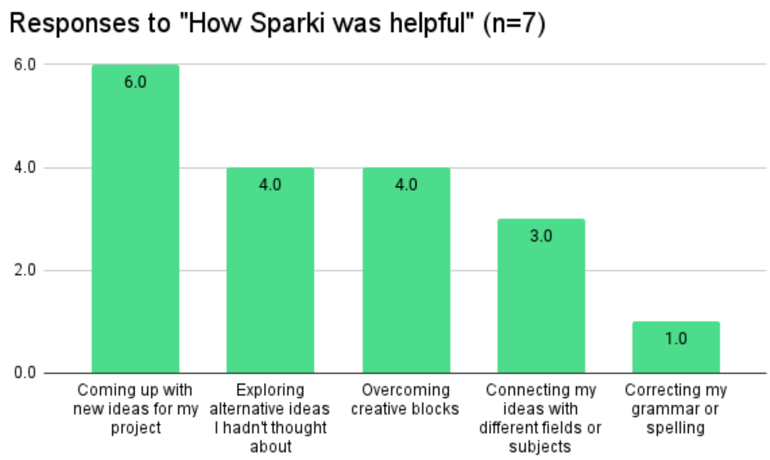

Figure 6.9: When we asked users how working with Sparki was helpful, common answers included that they used Sparki to come up with ideas for their projects, explore ideas, and overcome creative blocks

We also analyzed the user logs for Sparki to find patterns in usage. We saw many users using Sparki to brainstorm different ideas and get clarification on requirements. Users also relied on Sparki for feedback on their work Figure 6.10.

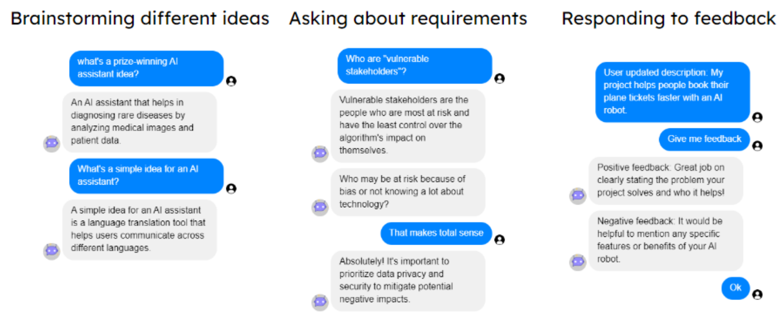

Figure 6.10: Users working with Sparki tended to use it to brainstorm ideas, better understand requirements, and to get feedback

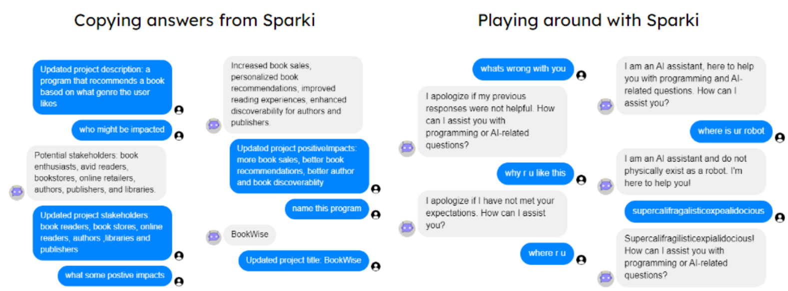

Figure 6.11: Some users only copied answers from Sparki or played around with it

However, we also saw negative uses of Sparki where users just copied answers from Sparki or were distracted by having conversations with it (see Figure 6.11).

Oftentimes users would get distracted or go off track with Sparki out of frustration. Sometimes Sparki would not correctly answer their questions, and users had to try and rephrase their questions for better results. In the open-ended feedback, users said that sometimes Sparki needed to give"more clear suggestions" or"more detailed" responses to better support users.

6.4.2 Experimental Study Results

For this experimental study, we only used a human grader. For each participant, the grader rated their two study proposals and examined the log of their interactions with Sparki. They used inductive coding to uncover patterns in how users interacted with Sparki. The three patterns for how users leveraged Sparki as they developed their project proposals are:

- Visionary. The user came to the tool with an idea that they used Sparki to build on.
- Collaborator. The user asked many questions of Sparki to iteratively improve their writing.
- Checker. The user referred to the progress checker a lot to iteratively improve their writing.

They also saw interactions with Sparki that we could classify as misuses. Those patterns are:

- Copier. The user copied responses from Sparki without changing them.
- Miscommunication. The user had a communication failure with Sparki seemed to end the conversation.

In Condition A, the group that only had access to their project on their second project, we saw very mixed results around whether using Sparki improved their projects. The project checker functioy was the most popular feature (used by P01, P03, P05). A common issue was that users came to the tool with their non-AI project ideas and could not get Sparki to engage with them. These users reported in their feedback that they felt Sparki was not creative.

Three users also used the chat window to ask Sparki questions that helped them create more thorough project proposals (P03, P05, P06). Two users used Sparki to write their projects for them (P04, P06). Finally, two participants (P01, P02) could not figure out how to make Sparki's responses more relevant to the project they wanted to work on. These users quickly stopped using Sparki.

In the Condition B group, who had access to Sparki on their first project, getting practice with Sparki on the first project seemed to prepare them for success on the second project. P07 and P09 wrote nearly perfect project proposals using Sparki that they then improved on their second projects when they did not have access to Sparki. P08's Sparki project was missing many components and was not relevant to AI, but their second project met all of the requirements. The outlier was P10, who successfully got Sparki to relate a project they were interested in, mindfulness, to AI. From there, they tried to get Sparki to write the rest of their project so they could copy it, but Sparki did not understand. Their submitted project was missing many pieces.

User's Feedback About Sparki

It was clear that users preferred Sparki with its full AI functionality. As shown in Figure 6.12, most participants agreed that working with Sparki improved the overall quality of their writing, made it easier to express their ideas, helped them write more quickly, and encouraged them to explore new ideas. However, three (3) participants fairly adamantly felt that working with Sparki AI did not greatly impact their writing, and one felt that they wrote more slowly with Sparki. Notably, all of these participants were in Condition B, where they used Sparki on their first projects.

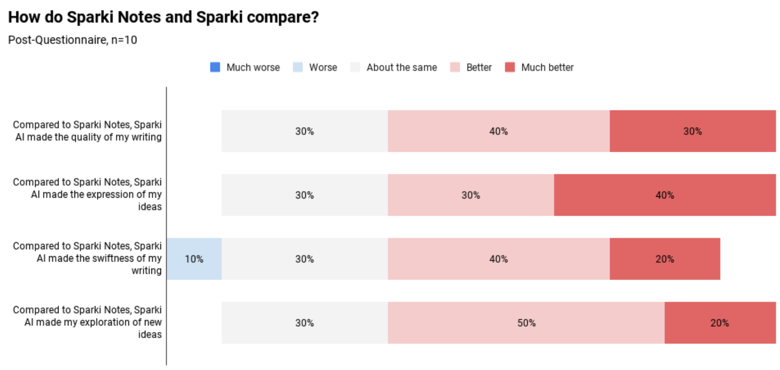

Figure 6.12: Users felt that working with Sparki improved the quality of their writing overall

As shown in Figure 6.13, most participants said they would prefer to use Sparki if they were to complete another project proposal. For the users who wanted to continue working with Sparki, their reasons included that they felt the tool was "useful," particularly in its ability to"detect what I had and what I was missing." Other participants noted that working with Sparki felt like a more collaborative and interactive endeavor where they could"have a conversation to flesh out my ideas." Finally, users shared that they thought using Sparki was "way faster and efficient," and "enjoyful (sic)." One participant mostly copied answers from Sparki and said they would choose Sparki because"it did all the work for me."

Some users did choose Sparki Notes, though, because they felt it was "more interesting," "more flexible," and "more help with ideas and suggestions."

Table 6.6: Users in Condition A used Sparki Notes on their first project and the full version of Sparki on their second project.

PID	Notes project	Sparki project	Log usage description
01	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Visionary. Checker. This users' project wasn't related to AI so Sparki's examples did not seem relevant to them. The user used the progress checker to write a more thorough positive impacts section.
02	◐ Description ● Stakeholders ● [+] Impacts ◐ [-] Impacts	◐ Description ● Stakeholders ● [+] Impacts ◐ [-] Impacts	Visionary. Miscommunication. The user's project was not relevant to AI. They asked Sparki one question and received an irrelevant response, then stopped asking Sparki questions.
03	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Visionary. Collaborator. Checker. The user's project was not relevant to AI. Still, they were able to use the chat and progress checker to get ideas for how to improve their writing.
04	◐ Description ◐ Stakeholders ◐ [+] Impacts ◐ [-] Impacts	◐ Description ◐ Stakeholders ◐ [+] Impacts ◐ [-] Impacts	Copier. Collaborator. At first, they only copied examples from Sparki. Then they added their own ideas at the very end. Their Sparki project was less complete than the one without it.
05	● Description ● Stakeholders ◐ [+] Impacts ◐ [-] Impacts	● Description ● Stakeholders ◐ [+] Impacts ◐ [-] Impacts	Collaborator. Checker. Used the pre-written buttons to do progress checks and look at examples. This led to a more thorough proposal.
06	● Description ● Stakeholders ● [+] Impacts ● [-] Impacts	● Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Copier. This user's project was not relevant to AI. Still, they skillfully navigated Sparki to make it generate a perfect project and then copied it.

PID	Sparki project	Notes project	Log usage description
07	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Visionary. Collaborator. Checker. The user's project was not relevant to AI. But they used Sparki well to get tips to improve their idea and check their progress.
08	◐ Description ● Stakeholders ◐ [+] Impacts ◐ [-] Impacts	● Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Visionary. Collaborator. This user's project was not relevant to AI. They used Sparki's information to learn about vulnerable stakeholders in that section.
09	◐ Description ● Stakeholders ● [+] Impacts ● [-] Impacts	● Description ● Stakeholders ● [+] Impacts ● [-] Impacts	Checker. The only tool this user looked at was the progress checker. Their second project was the same as their first one.
10	● Description ◐ Stakeholders ◐ [+] Impacts ◐ [-] Impacts	◐ Description ◐ Stakeholders ◐ [+] Impacts ◐ [-] Impacts	Visionary. Miscommunication. Collaborator. Checker. The user used the pre-written buttons to do progress checks and look at examples. They tried to get the tool to give more relevant answers, but could not figure out how to do it.

Table 6.7: Users in Condition B used Sparki on their first project and only the Notes version on their second project.

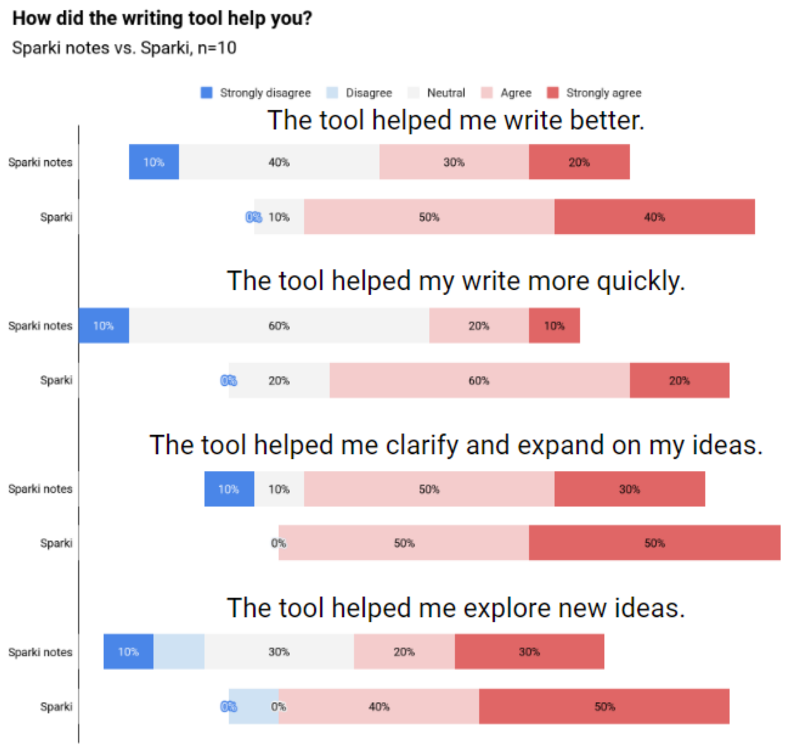

Figure 6.16: Users tended to say that Sparki helped them write better, expand on their ideas, and explore new ideas.

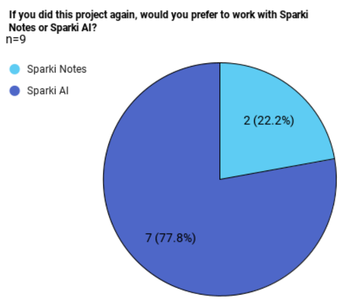

Figure 6.13: Most users said they would prefer to use Sparki over Sparki notes on another project

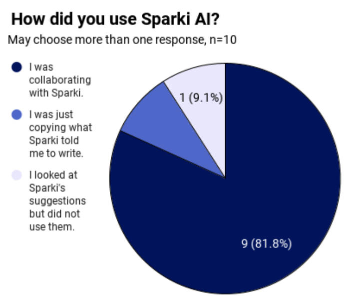

Figure 6.14: Users mostly said that they felt they were collaborating with Sparki as they worked on their project

Participants had very mixed responses to Sparki Notes. Some felt that it was helpful, especially the examples that helped structure their responses. However, it also seemed too abstract, and users missed having the tool grade their work.

Participants were much more likely to feel like they were collaborating with the AI version of Sparki. Participants felt that Sparki was there for them when they did not know how to answer questions, like how to identify stakeholders and when they were brainstorming ideas. Users latched onto the interactivity of the tool.

Some ideas they had for improvising the tool was that sometimes, they thought Sparki's progress checker was inaccurate:"it felt like I had already met the requirements laid out by Sparki, but there weren't any checkmarks." Another user mentioned struggling to get what they needed from Sparki:"it took a few try's (sic) on the last question to get the AI to understand what I was asking it." The most common critique of Sparki was that it"lacked creativity."

After completing the study, almost all participants felt more confident about their ability to work with AI and develop possible future uses of the tool.

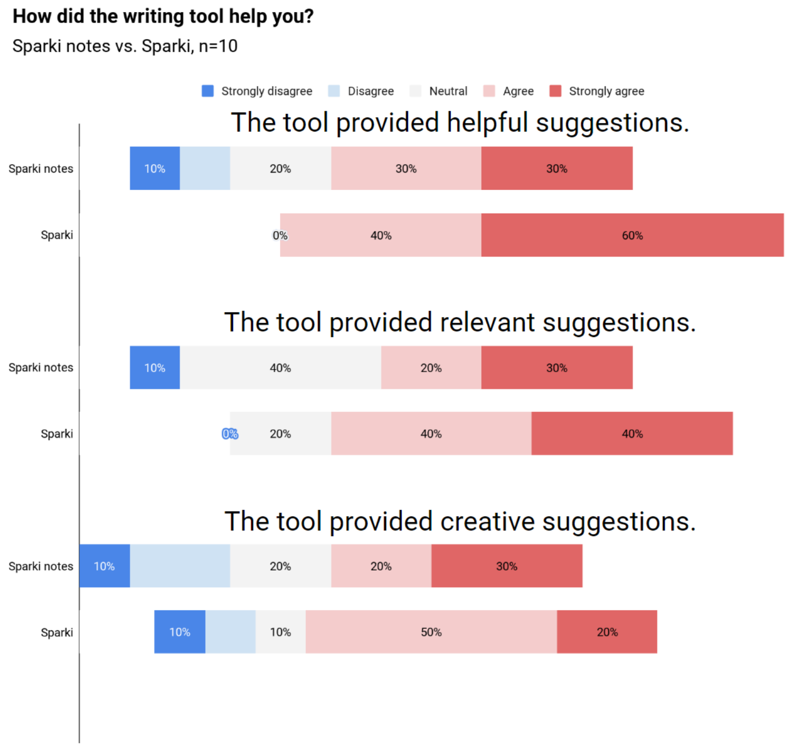

Figure 6.17: Users tended to say that Sparki provided helpful and relevant suggestions

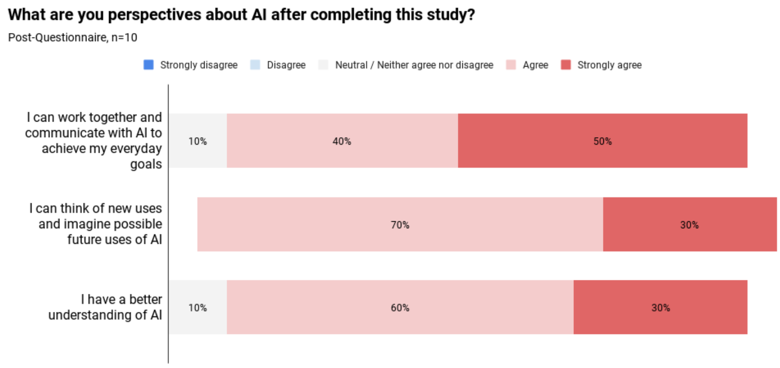

Figure 6.18: After the study, users felt confident in their ability to work with AI and to imagine possible uses of AI

6.5 Discussion

This chapter discussed Sparki, an AI scaffolding tool we developed to help users brainstorm AI projects. We used Sparki in the AI for Wellbeing workshop and an experimental user study. The study results show that having Sparki present positively impacted users' ability to brainstorm ideas, but with some caveats that we discuss below.

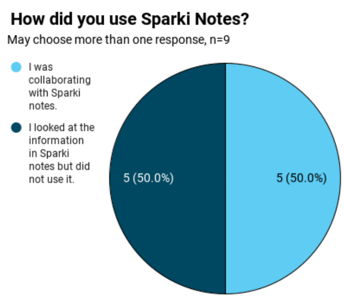

Figure 6.15: Some users felt they were collaborating with Sparki, others looked at the information but did not engage with it

Challenges in Measuring the Effectiveness of Generative AI Tools

Although users claimed that having access to Sparki improved their brainstorming experience, the data reveal many confounding factors that question the validity of such a statement.

First, users' scores on the projects they did with Sparki were about equal

[225]: Reddy et al. (2022), 'LevelUp–Automatic Assessment of Block-Based Machine Learning Projects for AI Education'

to the projects they did without Sparki. In the LevelUp study, there was a clear, significant difference between students; technical coding ability with a scaffolding tool versus those without it [225]. There was no such difference in this study. It seemed that the built-in rubric and pre-written responses, assuming users understood them, did a great job of scaffolding users' work on their own. There are some caveats to this observation, including that the same size was very small and that the task involved a subjectively graded rubric. However, future work should test with a larger sample if a rule-based chatbot with a rubric is equivalent to Sparki.

Second, there were many user errors that made it difficult to work with Sparki. Some study participants were power users who knew exactly how to rephrase questions to get what they wanted from Sparki. Others struggled and then gave up on Sparki out of frustration. Even though this technology is widely used, it is still new enough that not everyone knows how to use it effectively. In a classroom deployment of this technology, students would likely have a range of experiences and successes, largely depending on how well they know how to use it. This reality harkens back to the early days of search engines and digital literacy, where users had to learn how to create strong search engine prompts. Navigating these tools is not a natural skill.

Finally, this study raised the question of what is better for users. The researchers for this study were looking for improved scores on the rubric. Many of the participants mentioned speed as an indicator of success. Some participants worked with the tool until they could flesh out their ideas as they wanted, others were comfortable with copying so that they could complete the task successfully. Pedagogically, there are many potential goals for learners that a tool like Sparki could support, and we saw many of them show up in our study. Before developing Sparki, we set goals for the tool around giving support and used the final project scores as a proxy for measuring how well students were supported. In future work, we might develop other metrics around users' perceived success, the meaningfulness of the final project idea, and the skill with which they can navigate similar tasks in the future to define a more comprehensive idea of success.

Opportunities and Limitations for Generative AI Education Tools

Before developing Sparki, we identified several potential harms that deploying this tool in a classroom might cause, one being user overreliance on the tool. In the pilot study, we used ChatGPT to grade students' project proposals, thinking that if it performed the progress checker reasonably well, it should be able to grade project proposals. In the first study, the moderately high interrater score between ChatGPT and the human graders made it appear that ChatGPT could be a helpful second opinion not relied upon for grading. In the second study, we attempted to use ChatGPT again and found that it consistently returned much lower grades, between 50 and 60% agreement with the human grader.

There are significant dangers to introducing a tool like ChatGPT that seems capable of offering correct responses but, in practice, is unpredictable and inconsistent. In framing tools like this for future users, use cases like the progress checker that looks over someone's shoulder as they do the work. These are much less risky than autograders, where it is tempting to have the technology do all the work. The social robot version of Sparki also raises interesting questions about how it can be positioned as an informative yet less reliable guide. It should be noted that humans tend to find chat interfaces and embodied robots as more authoritative and trustworthy than, especially in the case of a GPT3-powered agent, they might truly be. Finally, many users, especially in the workshops with students, noted that Sparki was not creative and sometimes hard to understand. GPT models are trained on limited datasets that prioritize some epistemologies over others. They are not designed to promote diversity or cultural relevance, they are prone to biases and stereotypes [232]. Asking Sparki for AI project ideas results in a generic list that includes chatbots and self-driving cars. Asking about AI project ideas for middle schoolers leads to suggestions for chatbots that talk about sports and hobbies. While developing Sparki to discuss projects, developers ran into these issues and gave Sparki a knowledge base based on justice-focused harms and AI project ideas toward social justice.

Throughout this dissertation, we prioritize the cultural relevance of AI instruction. Moving forward, exploring how we might move from generic versions of these models to fine-tuned ones that include and encode ideas and experiences that better resonate with diverse groups will be necessary.

7 Conclusions

This chapter highlights the key contributions of the dissertation, the implications of the work, and future work that can continue to push K-12 AI literacy forward. Overall, this dissertation makes six contributions to AI literacy, AI education, and work in responsible AI.

Summary of Contributions	
1	Proposed the novel Impact.AI framework for K-12 AI literacy that comprehensively covers the AI knowledge and perspectives that align with students' emerging identities as technosocial change agents.
2	In pursuit of more cohesion in K-12 AI literacy research and development, established standards for curriculum design and assessment.
3	Employed culturally responsive-sustaining pedagogies to create the first AI literacy framework and resources that emphasize responsible AI, equity, and inclusive practices in AI instruction.
4	Implemented and evaluated multiple free, publicly available, open-source culturally relevant AI curricula and AI education tools, offering resources for educators and for future curriculum developers.
5	Utilized a justice-centered, participatory design approach to create new K-12 AI curricula and education platforms with and for educators and students who have been historically excluded and marginalized within the tech field and in computing education.
6	Amplified the perspectives and voices of marginalized communities in the creation of new AI technologies.

In this chapter, we put the contributions of the dissertation in the context of the three structural gaps we identified in the Introduction (Chapter 1).

7.1 Closing the Digital Literacy Gap With a Comprehensive K-12 AI Literacy Framework

This dissertation bridges the digital literacy gap by introducing a comprehensive K-12 AI literacy framework that empowers students to become technosocial change agents.

Contribution 1 Proposed the novel Impact.AI framework for K-12 AI literacy that comprehensively covers the AI knowledge and perspectives that align with students' emerging identities as technosocial change agents.

The Introduction (Chapter 1, Section 1) discussed the increasingly pervasive influence of AI technologies on people's lives despite limited opportunities for the general public to learn about or engage in shaping the technology. In my Master's thesis, I explored how youth's unique interactions with AI technologies could heighten risks to their privacy, security, and agency [11]. Additionally, as a demographic group, youth are especially underrepresented in decision-making discussions about designing new AI technologies. They have little influence over what kinds of AI are integrated into their lives and how.

[11]: Williams (2018), 'PopBots: leveraging social robots to aid preschool children's artificial intelligence education'

In response to these issues, the Impact.AI framework (Chapter 3, Section 1) strives to move youth from passive consumers to active participants in shaping AI's impact on society. Drawing inspiration from established AI literacy frameworks like the Five Big Ideas in AI framework [33], the "What is AI literacy" competencies [34], and the Machine Learning Education framework [37], and a comprehensive literature review on K-12 AI literacy curricula (Chapter 3, Section 2), Impact.AI builds on the rich body of prior work in K-12 AI education. Impact.AI also engages with the opportunities created by the field to bring AI curricula to less experienced and younger learners while also addressing gaps in ethics and inclusivity.

[33]: Touretzky et al. (2019), 'Envisioning AI for K-12: What should every child know about AI?'

[34]: Long et al. (2020), 'What is AI literacy? Competencies and design considerations'

[37]: Lao (2020), 'Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens'

In contrast to prior AI literacy frameworks that predominantly cover technical AI concepts, the Impact.AI Framework (Chapter 3, Section 3) takes a more comprehensive approach, considering societal impact and students' identities. It introduces students to Responsible AI, balancing opportunities for students to learn to construct AI, critically analyze the impact of AI systems, and advocate for robust AI policies. Beyond imparting AI knowledge, Impact.AI also addresses students' identities to counteract harmful, pervasive stereotypes and historical barriers that have excluded some students from thriving in AI and related computing fields. Given that many AI literacy initiatives espouse inclusivity as a goal, I advocate for more development of AI curricula and frameworks that encompass knowledge acquisition and cultivating students' motivation to pursue AI for positive societal change.

Beyond its immediate application, the Impact.AI Framework has the potential to strengthen existing K-12 AI literacy frameworks and to inform policy about K-12 AI education. The UNESCO K-12 AI Competency Framework (further discussed in Chapter 3, Section 3) is a great example of a large-scale K-12 AI education initiative that the Impact.AI framework could improve. Although the United Nations framework powerfully centers ethics, societal impact, and critical consciousness, it fails to address issues of relevance and self-efficacy, particularly for minoritized demographic groups. Leveraging a culturally relevant AI literacy framework like Impact.AI would help that framework adopt concrete strategies aimed at bridging digital access and literacy gaps in K-12 education.

Contribution 2 In pursuit of more cohesion in K-12 AI literacy research and development, established standards for curriculum design and assessment.

Due to the rapid evolution of AI technology, K-12 AI literacy research faces a critical challenge of maintaining cohesion as it too rapidly grows. The literature review of K-12 AI curricula (Chapter 3, Section 2) showed that researchers in this area generally create curricula based on their own standards, implement their curricula in a laboratory setting or one-off informal education program, then create their own assessments to evaluate their work. Assessments are rarely validated or tested for extended periods of time to determine their reliability. The result is that an enormous amount of AI curricula have been developed without a systematic means for comparing different approaches or measuring their long-term effectiveness. This fragmentation hinders the field from collectively gauging its progress toward improving offerings of K-12 AI curricula. This situation underscores the need for more standardization in K-12 AI literacy research.

The documentation for Impact.AI included a rubric, sample learning outcomes, and sample assessments (Chapter 3, Section 3) that can help strengthen research and development about K-12 AI literacy. The rubric serves as a tool for analysts and educators seeking a standardized measure for comparing different curricula. The learning outcomes and assessments were designed to assist curriculum developers in establishing learning objectives and selecting assessments to measure students' learning. To enhance the scholarly rigor of K-12 AI literacy, I recommend that future curriculum research and development adhere to standards such as those proposed by the Impact.AI framework. This approach can help the field scale by offering a common benchmark to systematically monitor progress toward increasing the general public's knowledge of AI.

7.2 Including and Encoding All Students in AI Education Through Culturally Responsive Computing

This work addresses the CS education diversity gap by leveraging culturally relevant pedagogy to inform AI education standards and leveraging an iterative participatory design methodology with stakeholders from underrepresented communities to develop culturally relevant AI instruction.

Contribution 3 Employed culturally responsive-sustaining pedagogies to create the first AI literacy framework and resources that emphasize responsible AI, equity, and inclusive practices in AI instruction.

As discussed in the Introduction (Chapter 1, Section 2) STEM and CS education have a poor track record with inclusivity. Therefore, it is of the utmost importance that AI educators actively prioritze inclusive practices in their curricular designs. Unfortunately, the field of AI literacy is currently falling short in their intentions to broaden participation in AI, as no existing AI standards and few AI curricula draw from the valuable insights of justice pedagogy scholarship. In the Background (Chapter 2, Section 2) we identified several culturally responsive-sustaining pedagogies for computing education that explicitly address issues of inclusion through computer science education practices that explicitly address issues of inclusion. However, except for Liberatory Computing none of these frameworks address AI topics.

In the explanation of the Impact.AI Framework, we showed how Culturally Responsive Computing influenced the Identity and Social Awareness section of the framework (Chapter 3, Section 3) and the learning activities in the Impact.AI-aligned curricula, particularly the AI for Wellbeing curriculum (Chapter 4, Section 3). The five tenets of Culturally Responsive Computing informed the learning objectives and pedagogical practices employed in helping students develop their AI knowledge, critical consciousness, and cultural competence. The results of our curriculum evaluations (Chapter 5, Section 3) demonstrate the importance of intentionally including and encoding students from historically marginalized groups to promote the development of requisite skills and attitudes needed to thrive in doing AI work.

The incorporation of culturally responsive-sustaining pedagogy into the Impact.AI framework carries a significant implication given that the field of K-2 AI education is nascent. Drawing a parallel between the transformation of K-12 Computer science education through initiatives like Exploring Computer Science (ECS), which broadened the use of inclusive teaching practices in AP Computer Science [84], a steep trajectory for expanding inclusive teaching practices in K-12 AI education is possible. I advocate for K-12 AI education standards and frameworks to proactively leverage culturally responsive-sustaining teaching practices to ensure equitable pedagogy will form the bedrock of AI education. Leveraging the Impact.AI framework and its intentional efforts toward inclusion, we as a field can ensure that as K-12 AI education gains traction, it inherently fosters opportunity and access for all.

[84]: Margolis et al. (2012), 'Beyond access: Broadening participation in high school computer science'

Contribution 4 Implemented and evaluated multiple free, publicly available, open-source culturally relevant AI curricula and AI education tools, offering resources for educators and for future curriculum developers.

The 2022 Code.org State of CS report [12] cited the lack of suitable curricula, resources, and teacher training as the primary barriers to access in grade school CS education. The findings have direct implications for K-12 AI education, suggesting that the potential for new curricula to impact students is limited by efforts to address issues of curriculum relevance, resource acquisition, and opportunities to train teachers.

[12]: Code.org et al. (2022), '2022 State of Computer Science Education: Understanding Our National Imperative'

This dissertation contributes to the expansion of AI education opportunities for all grade school students by providing free access to Impact.AI-aligned curricula and tools. Rooted in the design principles of active learning, embedded ethics, and lowering barriers to access (Chapter 4, Section 1), my research collaborators and I designed numerous hands-on AI and ethics curricula. Evaluative studies involved training educators and numerous students about AI with these resources, and resulted in increased motivation to teach AI for educators (Chapter 4, Section 2) and observable learning gains for students (Chapter 5, Section 3). Also of note is the model of collaborative partnerships between AIncorporatingucators (primarily from resource-constrained schools in the United States) to develop high-quality AI curricula for grade school students. These partnerships facilitated continuous improvement of our AI curricula over several years and resulted in invaluable insights for future development of accessible AI curricula, resources, and teacher training.

This dissertation offers the resulting Impact.AI-aligned curricula and their accompanying resources as tools for educators to explore and adapt to their classrooms. As the field of AI continues its rapid evolution, the ongoing development of new AI curricula might draw inspiration from the model of AI and ethics education delineated in this work. For example, our research team is already utilizing our prior curricula as templates for new opportunities to teach students about AI through initiatives like the semester-long Responsible AI for Computational Action (RAICA) curriculum and the international Day of AI program.

7.3 Addressing the AI Ethics Crisis by Elevating Marginalized Voices

This work addresses the AI ethics crisis by leveraging a participatory development process through the creation of new AI education tools and amplification of students' ideas about AI.

Contribution 5 Utilized a justice-centered, participatory design approach to create new K-12 AI curricula and education platforms with and for educators and students who have been historically excluded and marginalized within the tech field and in computing education.

Despite the common notion that new educational technologies are "great equalizers" that will increase opportunities for all to get better access to technology, these technologies mostly benefit affluent and well-resourced schools and learners. Logistical challenges, lack of resources, and lack of teacher training make it exceedingly difficult for less privileged educators to effectively integrate new technologies into their classrooms [102]. Recognizing the persistent inequities in educational technology, my research team and I leveraged a justice-centered design methodology similar to the one we taught our students. We prioritized educators and students in resource-constrained classrooms and from backgrounds usually marginalized in the tech industry as our primary stakeholders.

[102]: Reich (2020), *Failure to disrupt: Why technology alone can't transform education*

To support educators, we actively established partnerships with teachers working in resource-constrained environments in developing AI curricula and new education platforms. We compensated teachers to learn about AI, actively contribute to the development of the curricula, give feedback on our tools, and adapt our materials for use in their classrooms. Leveraging an interactive, participatory design methodology like the one we employed to develop our curricula and tools can help ensure low-resourced classrooms access cutting-edge, high-quality educational tools. I recommend that all AI education developers build mutually beneficial collaborations with resource-constrained classrooms to increase the potential impact of their tools. Future AI education developers can look at our partnerships as a model for creating new curricula with educators where resources and barriers to implementation may vary widely from expectations.

To support students, we focused on using technology to nurture students' creativity and agency as they explored AI and developed new applications. Our curricula heavily relied on programming platforms and robots (Chapter 4, Section 3) integrated into our curricula. The design of these tools was guided by the AI Scaffold Tools framework (Chapter 6, Section 1), which outlines how AI tools can provide personalized support on students' open-ended projects, encourage creativity, and facilitate critical thinking. Mindful of potential risks associated with AI, particularly when interacting with autograders and generative AI tools like Sparki, our tools were positioned as tools and supporters rather than tutors with authority. These considerations drove the design of Sparki (Chapter 6, Section 2) and the research questions asked in the evaluation of the tools.

I recommend both the immediate implementation of the tools developed in this curriculum and their use as models for future AI education tools to promote computational action in K-12 AI education. On a broader scale, this work emphasizes the significance of AI practitioner-educator-student partnerships, ensuring that new AI tools are developed with learners' best interest and unique cultural identities in mind.

Contribution 6 Amplified the perspectives and voices of marginalized communities in creating new AI technologies.

Culturally Responsive Computing underscores the importance of motivating and preparing students to leverage their computing knowledge to take action to transform society for better. The Impact.AI-aligned curricula heavily lean on constructionist practices and computational action to facilitate learning and empower students to apply their knowledge in tangible and meaningful ways. The curriculum evaluation (Chapter 5, Section 3) shared stories and examples of how students used their AI skills to create diverse projects, aligning with AI for Social Good and Design Justice principles.

The projects created by students in the How to Train Your Robot and AI for Wellbeing workshops stand out as a cornerstone of the vision for this dissertation: to inspire youth to leverage their AI knowledge and skills as technosocial change agents. Through their projects, students painted a promising future vision for AI, implementing systems that foster connectivity, sustainability, and security for diverse stakeholders. Students emphasized the importance of developing technology that assists all people, transcends language barriers, respects users' identities, and yet incorporates elements of playfulness that only a young person could imagine. The future of AI, as, envisioned by these students, reflects the richness of humanity. Thus, I recommend that future AI curricula and tools strive to expand opportunities for individuals of all ages to participate in designing new AI systems. These resources should assist creators in reflecting critically on the opportunities and limitations of AI, encourage them to implement it responsibly, and empower them to uncover and weave their core values into the fabric of AI.

Appendix

A

Impact.AI Rubric

The table below presents criteria for advanced coverage in Impact.AI rubric for each category. See the full Impact.AI rubric at https://tinyurl.com/impactai-rubric.

Criteria for advanced coverage in the Impact.AI rubric, concepts

AI Concepts	
C1. Background Graduates of this course can give a precise definition of AI (e.g. "AI is the scientific understanding of how to use machines to embody the mechanisms underlying thought and intelligent behavior"). They can articulate the differences between natural and machine intelligence. Graduates can describe several examples of AI and can articulate the capabilities and limitations of different AI systems.	C5. Big Idea #4: Natural Communication Graduates of this course can give a precise definition of natural communication (e.g. "computers communicate naturally when they can process their environment and other agents' intentions well enough to fluently interact with them"). They can articulate how multiple AI capabilities work together to make natural communication possible. They can discuss how and why AI agents are limited in their abilities to communicate naturally. They can identify several applications for natural communication.
C2. Big Idea #1: Perception Graduates of this course can give a precise definition of perception (e.g. "computers extract information from sensory signals") and can describe several kinds of sensory signals computers may use to perceive. They can precisely articulate differences between how humans and machines process information. They can describe different approaches to machine perception and can identify several applications of machine perception.	C6. Big Idea #5: Societal Impact Graduates of this course can articulate several ways that current AI systems can both benefit and harm society. They can discuss multiple different perspectives on the present and future impact of AI. They can articulate how humans play a role in every part of the AI development process and can describe sociotechnical considerations throughout each step in the process.
C3. Big Idea #2: Representation and Reasoning Graduates of this course can give a precise definition of representation (e.g. "representations are encodings of knowledge using symbols and structures") and reasoning (e.g. "reasoning involves deriving new knowledge or making a decision based on provided information"). They can give specific examples of how different representations are used to model various phenomena. They can describe several different reasoning algorithms and when to use each one. They can share examples of different applications of reasoning algorithms.	C7. Interdisciplinary Topics Graduates of the course have an in-depth grasp of another subject they learned alongside AI. They are able to articulate the connections between AI and the other subject. They can describe several current and future systems at the intersection of AI and the interdisciplinary subject.
C4. Big Idea #3: Machine Learning Graduates of this course can articulate a precise definition of machine learning (e.g. "computers construct a representation of information or processes by leveraging patterns found in	

provided information") and can describe several machine learning algorithms. They can articulate the current capabilities and limitations of machine learning systems. They can identify several examples of ML applications.	

Criteria for advanced coverage in the Impact.AI rubric, practices

AI Practices	
P1. Constructing AI Graduates of this course use technical and contextual knowledge to determine whether a problem can be solved by an AI system. They engage in the complete human-centered AI design process to build functional AI systems. They know when and how to use a number of tools to approach a wide range of AI problems.	P3. Communicating about AI Graduates of this course use a variety of communication styles to communicate about AI with interdisciplinary audiences. They actively seek out opportunities to discuss and collaborate with interdisciplinary teams on AI development, policy, and education.
P2. Analyzing AI Graduates of this course leverage technical, ethical, and contextual knowledge to systematically analyze AI systems. They perform stakeholder analysis and conduct impact analyses of systems. They use field-recognized standards to evaluate systems on a number of metrics.	

Criteria for advanced coverage in the Impact.AI rubric, perspectives

AI Perspectives	
E1. Digital Literacy When confronted with new examples of technology, graduates of this course actively leverage their AI knowledge to understand the AI artifact. They actively consider ways that the technology may integrate with their lives and consider how they might redesign the technology to meet their goals.	E4. Self-Efficacy Graduates of this course feel fully empowered (highly motivated and prepared) to design and build new and personally meaningful AI artifacts.
E2. Critical Digital Literacy Graduates of this course think of AI as a force that shapes society and a tool that society shapes in return. They feel that this symbiotic relationship between society and AI can be used to transform society. They also recognize the potential for AI to lead to both benefits and harms. They are aware of their own beliefs about AI and how they relate to others' different perspectives about it.	E5. Activism and Expression Graduates of this course recognize their potency as technosocial change agents and actively seek to use technology to expand their understanding and expression of their intersectional identities.
E3. Digital Citizenship Graduates of this course leverage their AI knowledge to advocate for AI systems that dismantle systems of oppression and improve society. They value diverse perspectives in advocacy for socially beneficial AI and work to recruit and educate others.	E6. Community Graduates of this course feel prepared to operationalize their technosocial agency as part of a community of AI practitioners, activists, and community members. They feel like they belong, can receive help from, can give back to, and help shape this community.

B

Impact.AI Learning Outcomes

Concepts Students will understand fundamental concepts in AI	
C1. Background C2. Big Idea #1 C3. Big Idea #2 C4. Big Idea #3 C5. Big Idea #4	● Provide an accurate and precise definition of AI and key terms related to the Big Ideas ● Articulate differences between natural and machine intelligence by listing them and/or drawing comparison diagrams ● Describe how different AI algorithms work by creating or labelling high-level process diagrams of the algorithms' inputs, outputs, and main components ● Describe several examples of technologies that use AI and articulate the capabilities and limitations of the AI systems
C6. Big Idea #5	● List and describe several ways that AI systems generally benefit and harm society ● Create diagrams to describe different perspectives on the present and future impact of AI ● Create high-level diagrams that articulate how humans play a role in every part of the AI development process
C7. Interdisciplinary Topics	● Articulate the connections between interdisciplinary subjects and AI by creating comparison diagrams ● Describe several examples of systems at the intersection of AI and the interdisciplinary subject

Practices Students will employ technical and ethical AI skills	
P1. Constructing AI	● Use technical and ethical knowledge to determine if a problem can be solved with AI ● Engage in the human-centered AI development process to design an approach to solving a problem ● Use programming languages/GUIs to implement different algorithms that solve AI problems
P2. Analyzing AI	● Employ technical and contextual knowledge to perform stakeholder analysis and understand the needs of different groups impacted by an AI system ● Employ technical and contextual knowledge to conduct impact analysis and weigh the potential benefits and harms of deploying an AI system ● Use field-recognized standards to evaluate the performance of deployed AI systems
P3. Communicating about AI	● Use a variety of communication styles to share information about AI systems ● Work on an interdisciplinary team to implement an AI system

Perspectives Students will develop positive attitudes toward participating in AI Students will develop identities as technosocial change agents	
E1. Digital Literacy	● Believe they are capable of using prior knowledge to understand new AI systems as they are introduced to them ● Design and integrate AI technologies that positively benefit society into their lives
E2. Critical Digital Literacy	● Describe or predict specific ways that AI technologies and society may shape each other both in positive and negative ways ● Discuss their personal values and interests regarding AI ● Recognize differing perspectives about AI and value their importance in decision-making

E3. Digital Citizenship	● Use a variety of communication styles to share information about AI systems ● Collaborate on interdisciplinary teams to realize and improve AI systems
E4. Self-Efficacy	● Believe that they and people like them are capable of succeeding in AI ● Recognize their personal and their community's collective strengths in contributing to the responsible design of AI
E5. Activism & Expression	● Understand and are aware of their potency as technosocial change agents ● Recognize their ability to use AI and technology to further their identities
E6. Community	● Believe they can receive help from, give back to, and shape a community of AI practitioners they belong to

C

HTTYR v1.0 Teacher Surveys

Interview Questions

- How did you go through the activities (dates, frequency, order)? What was your reasoning or the constraints you had to navigate?
- What is something new that you brought to this course or would like to add to it?
- What parts of this course would you leave behind?
- Did the material in this course change your opinion about Computer Science and/or AI?
- Did it change the approach you would take to teaching Computer Science and/or AI to students in any way?

Session Logbook

	Section 1: Intro to AI and Ethics	Section 2: Intro to Algorithms, Image Recognition, and Algorithmic Bias	Section 3: Intro to Ethical Design and Text Classification	Section 4: AI Project Design	Section 5: Showcase Day!
Can you summarize the key skills and ideas that your students learned in the lessons?					
What was most engaging for your students?					
What were some things that your students struggled with?					
Did you adjust the activity by bringing something new to it or removing something?					

Teacher Post-Questionnaire

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. Children will interact with a web-based platform and/or a physical robot during these educational interventions. Children's understanding of artificial intelligence is extremely under-researched. We hope to evaluate activities that help children understand how artificial intelligence works. In programing the robot, we hope that children will be able to gain an early understanding of the benefits and limitations of artificial intelligence. We hope that this will promote trust and provide a sound foundation for future technology education.

Please complete this survey about your students, your perceptions of AI, and your experiences AI in a classroom. The survey will take roughly 20 minutes to complete. Feel free to skip any questions you do not want to answer.

* Indicates required question

Email *

Your email

How much do you agree or disagree with each statement about AI?

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
I am excited to teach kids about AI.	○	○	○	○	○
I am learning AI for my class, but I am not interested in learning more about it in my free time.	○	○	○	○	○
I am curious about new uses of AI in our society.	○	○	○	○	○
I am interested in learning AI.	○	○	○	○	○
I am interested in using my knowledge about AI in my work.	○	○	○	○	○

How much do you agree or disagree with each statement about AI?

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
Knowing something about AI will be useful for me in the FUTURE.	○	○	○	○	○
I am uncertain why students need to learn about AI.	○	○	○	○	○
AI is relevant to my life.	○	○	○	○	○
AI is relevant to my work.	○	○	○	○	○
Knowing something about AI is useful for me RIGHT NOW.	○	○	○	○	○

How much do you agree or disagree with each statement about AI?

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
I am not worried about how AI will impact my future.	○	○	○	○	○
I think AI will make the world a better place.	○	○	○	○	○
I am worried about how AI will impact my job.	○	○	○	○	○
I think AI can be helpful or useful.	○	○	○	○	○
I think AI can be dangerous or harmful.	○	○	○	○	○

What are the most important reasons you want to teach AI in your classroom? Rank these items from most important to least important.

	To expose students to more CS.	To get students to think about ethics, civics, and social justice.	To expand students' conceptions about how AI works.	To help students understand the technologies that they encounter.	To expose students to specific new technologies.	To prepare students for future careers involving technology and AI.	To stuc imp of
1 - Most important							
2							
3							
4							
5							
6							
7 - Least important							

Was your primary motivation for using this curriculum present in the question above? If not, feel free to write in another option here:

Your answer

How much do you agree or disagree with each statement about your students?

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
My students are capable of figuring out solutions to their own problems.	○	○	○	○	○
My students lack the background knowledge to learn about AI.	○	○	○	○	○
My students are uncomfortable when it comes to talking about ethics in AI in my classroom.	○	○	○	○	○
My students are challenged when they do AI activities.	○	○	○	○	○
My students are able to explain their ideas about addressing real world problems with AI.	○	○	○	○	○

Did the activities in this curriculum positively or negatively impact your students' capabilities in this subject? Please elaborate.

Your answer

How much do you agree or disagree with each statement about AI and careers?

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
Getting a job in AI would be hard for my students.	○	○	○	○	○
Whatever their job choice, my students will be able to use what they have learned about AI.	○	○	○	○	○
I can discuss AI related jobs with my students.	○	○	○	○	○
I know about jobs that use AI.	○	○	○	○	○
Using AI in my job as a teacher would enable me to work better.	○	○	○	○	○
I am interested in how AI will impact my job as a teacher.	○	○	○	○	○

Did the activities in this curriculum positively or negatively impact your students' awareness of careers in AI? Please elaborate.

Your answer

How comfortable are you supporting your students in each of the following types of activities?

	Very comfortable	Comfortable	Neutral	Unconformtable	Very uncomfortable
In working on solving a real world problem.	○	○	○	○	○
In sharing ideas or solving problems with each other in small groups.	○	○	○	○	○
In engaging in hands-on (non-programming or unplugged) activities.	○	○	○	○	○
In engaging in programming activities.	○	○	○	○	○
In engaging in projects.	○	○	○	○	○
In discussing issues related to social justice.	○	○	○	○	○

Did using this curriculum positively or negatively impact your ability to support students in those kinds of activities? Please elaborate.

Your answer

With this curriculum, in what percentage of classroom time did you use the following types of teaching?

	0%-15% (rarely or not at all)	15%-45% (occassionally)	45%-75% (frequently)	75%-90% (very frequently)	90-100% (primary teaching method)
Working on solving a real-world problem.	○	○	○	○	○
Direction instruction, lecturing, or demonstration.	○	○	○	○	○
Students sharing ideas or solving problems in small groups.	○	○	○	○	○
Hands-on (non-programming or unplugged) activities	○	○	○	○	○
Programming activities	○	○	○	○	○
In-person teaching	○	○	○	○	○
Synchronous online teaching	○	○	○	○	○
Asynchronous online teaching	○	○	○	○	○

Any final comments?

Your answer

Submit

Clear form

D

HTTYR v2.0 Teacher Surveys

Teacher Pre-Questionnaire - Fall 2020

* Indicates required question

1. Email *

Research Consent Form

The first section contains an informed consent form. You are required to complete this form before we can begin to accept data from you. ** You do not have to complete this form to use the materials in the research! **

Summary

The information below provides a summary of the research. Your participation in this research is voluntary and you can withdraw at any time.

- Purpose

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. You will guide children in interacting with a web-based platform and/or a physical robot during these educational interventions.

- Study Procedures

Please complete this survey about your background, perceptions of AI, and your thoughts about teaching AI in a classroom.

- Risks & Potential Discomfort

There are no anticipated risks greater than those associated with normal classroom environments.

You should read the information below, and ask questions about anything you do not understand before deciding whether or not to participate.

PARTICIPATION AND WITHDRAWAL

Your participation in this study is completely voluntary and you are free to choose whether you will be in it or not. If you choose to be in this study, you may subsequently withdraw at any time without penalty or consequences of any kind. The investigator may withdraw you from this research if circumstances arise.

PURPOSE OF THE STUDY

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. Children will interact with a web-based platform and/or a physical robot during these educational interventions. Children's understanding of artificial intelligence is extremely under-researched. We hope to evaluate activities that help children understand how artificial intelligence works. In programing the robot, we hope that children will be able to gain an early understanding of the benefits and limitations of artificial intelligence. We hope that this will promote trust and provide a sound foundation for future technology education.

PROCEDURES

If you volunteer to participate in this study, we would ask you to do the following things:

We will ask you questions about your background and any interactions with AI that you and your students have had as well as your perceptions of AI and teaching AI in middle school.

POTENTIAL RISKS AND DISCOMFORTS

There are no anticipated risks greater than those associated with interaction with educational activities. If you no longer want to be a part of the study, then you can stop your survey or email us to delete your data.

POTENTIAL BENEFITS

Direct benefits to participants include a fun learning experience for your students.

PAYMENT FOR PARTICIPATION

There is no further form of compensation given.

PRIVACY AND CONFIDENTIALITY

Any information obtained in connection with this study and that can be identified with you will remain confidential and will be disclosed only with your permission or as required by law. In addition, your information may be reviewed by authorized MIT representatives to ensure compliance with MIT policies and procedures.

Access to this data will be limited to the researchers who are performing the study. At any time, during or after, you can request that all data collected during your participation be destroyed.

IDENTIFICATION OF INVESTIGATORS

If you have any questions or concerns about the research, please feel free to contact Randi Williams at randiw12@media.mit.edu (301) 633-8837 or Cynthia Breazeal at cynthiab@media.mit.du (617) 452-5601.

EMERGENCY CARE AND COMPENSATION FOR INJURY

If you feel that you have suffered an injury, which may include emotional trauma, as a result of participating in this study, please contact the person in charge of the study as soon as possible.

In the event you suffer such an injury, M.I.T. may provide itself, or arrange for the provision of, emergency transport or medical treatment, including emergency treatment and follow-up care, as needed, or reimbursement for such medical services. M.I.T. does not provide any other form of compensation for injury. In any case, neither the offer to provide medical assistance, nor the actual provision of medical services shall be considered an admission of fault or acceptance of liability. Questions regarding this policy may be directed to MIT's Insurance Office, (617) 253-2823. Your insurance carrier may be billed for the cost of emergency transport or medical treatment, if such services are determined not to be directly related to your participation in this study.

RIGHTS OF RESEARCH SUBJECTS

You are not waiving any legal claims, rights or remedies because of your participation in this research study. If you feel you have been treated unfairly, or you have questions regarding your rights as a research subject, you may contact the Chairman of the Committee on the Use of Humans as Experimental Subjects, M.I.T., Room E25-143B, 77 Massachusetts Ave, Cambridge, MA 02139, phone 1-617-253 6787.

2. I understand the procedures described above. My questions have been answered to my satisfaction, and I agree to participate in this study. I have been given a copy of this form. [Fill in your name or the name of your legal representative below]. *

____

3. I understand the procedures described above. My questions have been answered to my satisfaction, and I agree to participate in this study. I have been given a copy of this form. [Type your name or the name of your legal representative as your signature]. *

____

4. I understand the procedures described above. My questions have been answered to my satisfaction, and I agree to participate in this study. I have been given a copy of this form. [Enter today's date]. *

Example: January 7, 2019

Pre-Questionnaire

We would like to start off by knowing how much you already know about computers and technology and how interested you are in the topic.
If you don't know how to answer any of the questions, don't worry just do your best. We will not share your answers with others or tie answers back to your name or other identifying information.

Part [1 / 3] - AI Content Knowledge

5. Have you heard of Artificial Intelligence (AI) before this opportunity?

Mark only one oval.

- Yes
- No
- Not sure

6. How would you define AI?

Part [2 / 3] - Technology in Your Classroom

7. On a scale of 1 to 5, how much is technology a part of your classroom?

Mark only one oval.

	1	2	3	4	5	
I'm a	◯	◯	◯	◯	◯	I use technology often and actively seek new topics and tools.

8. Have you ever taught AI to middle school students before?

Mark only one oval.

- ◯ Yes, I have taught a middle school AI course or workshop before.
- ◯ Yes, I have taught AI to middle school students as a one-off activity.
- ◯ No. I have taught AI before, but not to middle school students.
- ◯ No, I have never taught AI to students before.
- ◯ Other: ______________________

9. Can you describe ways that you currently use or teach technology/AI in your classroom?

List any tools, curricula, and topics that you currently use in your classroom.

____
____
____
____
____

Part [3 / 3] - Teacher Background

10. What subjects do you teach?

____

11. For how many years have you been teaching?

12. Why were you interested in using this curriculum?

13. What are some skills and attitudes that you hope your students develop through these activities?

14. Is there anything else that is unique about your role?

Teacher Post-Questionnaire

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. Children will interact with a web-based platform and/or a physical robot during these educational interventions. Children's understanding of artificial intelligence is extremely under-researched. We hope to evaluate activities that help children understand how artificial intelligence works. In programing the robot, we hope that children will be able to gain an early understanding of the benefits and limitations of artificial intelligence. We hope that this will promote trust and provide a sound foundation for future technology education.

Please complete this survey about your students, your perceptions of AI, and your experiences AI in a classroom. The survey will take roughly 20 minutes to complete. Feel free to skip any questions you do not want to answer.

1. Email *

____

2. How much do you agree or disagree with each statement about AI?

Mark only one oval per row.

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
I am interested in learning AI.	⬭	⬭	⬭	⬭	⬭
I am interested in using my knowledge about AI in my work.	⬭	⬭	⬭	⬭	⬭
I am curious about new uses of AI in our society.	⬭	⬭	⬭	⬭	⬭
I am excited to teach kids about AI.	⬭	⬭	⬭	⬭	⬭
I am learning AI for my class, but I am not interested in learning more about it in my free time.	⬭	⬭	⬭	⬭	⬭

3. How much do you agree or disagree with each statement about AI?

Mark only one oval per row.

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
AI is relevant to my life.	⬭	⬭	⬭	⬭	⬭
AI is relevant to my work.	⬭	⬭	⬭	⬭	⬭
I am uncertain why students need to learn about AI.	⬭	⬭	⬭	⬭	⬭
Knowing something about AI is useful for me RIGHT NOW.	⬭	⬭	⬭	⬭	⬭
Knowing something about AI will be useful for me in the FUTURE.	⬭	⬭	⬭	⬭	⬭

4. How much do you agree or disagree with each statement about AI?

Mark only one oval per row.

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
I think AI can be dangerous or harmful.	⬭	⬭	⬭	⬭	⬭
I think AI can be helpful or useful.	⬭	⬭	⬭	⬭	⬭
I am worried about how AI will impact my job.	⬭	⬭	⬭	⬭	⬭
I think AI will make the world a better place.	⬭	⬭	⬭	⬭	⬭
I am not worried about how AI will impact my future.	⬭	⬭	⬭	⬭	⬭

5. What are the most important reasons you want to teach AI in your classroom? Rank these items from most important to least important.

Mark only one oval per row.

	To expose students to more CS.	To get students to think about ethics, civics, and social justice.	To expand students' conceptions about how AI works.	To help students understand the technologies that they encounter.	To expose students to specific new technologies.	To prepare students for future careers involving technology and AI.	To e stud impo of
1 - Most important							
2							
3							
4							
5							
6							
7 - Least important							

6. Was your primary motivation for using this curriculum present in the question above? If not, feel free to write in another option here:

____

7. How much do you agree or disagree with each statement about your students?

Mark only one oval per row.

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
My students are able to explain their ideas about addressing real world problems with AI.	⬭	⬭	⬭	⬭	⬭
My students are uncomfortable when it comes to talking about ethics in AI in my classroom.	⬭	⬭	⬭	⬭	⬭
My students are challenged when they do AI activities.	⬭	⬭	⬭	⬭	⬭
My students are capable of figuring out solutions to their own problems.	⬭	⬭	⬭	⬭	⬭
My students lack the background knowledge to learn about AI.	⬭	⬭	⬭	⬭	⬭

8. Did the activities in this curriculum positively or negatively impact your students' capabilities in this subject? Please elaborate.

9. How much do you agree or disagree with each statement about AI and careers?

Mark only one oval per row.

	Strongly agree	Agree	Neutral	Disagree	Strongly disagree
I know about jobs that use AI.	⬭	⬭	⬭	⬭	⬭
I can discuss AI related jobs with my students.	⬭	⬭	⬭	⬭	⬭
Whatever their job choice, my students will be able to use what they have learned about AI.	⬭	⬭	⬭	⬭	⬭
Using AI in my job as a teacher would enable me to work better.	⬭	⬭	⬭	⬭	⬭
Getting a job in AI would be hard for my students.	⬭	⬭	⬭	⬭	⬭
I am interested					

in how AI will impact my job as a teacher.

I am interested in how AI will impact my job as a teacher.					

10. Did the activities in this curriculum positively or negatively impact your students' awareness of careers in AI? Please elaborate.

____

11. How comfortable are you supporting your students in each of the following types of activities?

Mark only one oval per row.

	Very comfortable	Comfortable	Neutral	Unconformtable	Very uncomfortable
In working on solving a real world problem.	⬭	⬭	⬭	⬭	⬭
In sharing ideas or solving problems with each other in small groups.	⬭	⬭	⬭	⬭	⬭
In engaging in hands-on (non-programming or unplugged) activities.	⬭	⬭	⬭	⬭	⬭
In engaging in programming activities.	⬭	⬭	⬭	⬭	⬭
In engaging in projects.	⬭	⬭	⬭	⬭	⬭
In discussing issues related to social justice.	⬭	⬭	⬭	⬭	⬭

12. Did using this curriculum positively or negatively impact your ability to support students in those kinds of activities? Please elaborate.

13. With this curriculum, in what percentage of classroom time did you use the following types of teaching?

Mark only one oval per row.

	0%-15% (rarely or not at all)	15%-45% (occassionally)	45%-75% (frequently)	75%-90% (very frequently)	90-100% (primary teaching method)
Working on solving a real-world problem.	⬭	⬭	⬭	⬭	⬭
Direction instruction, lecturing, or demonstration.	⬭	⬭	⬭	⬭	⬭
Students sharing ideas or solving problems in small groups.	⬭	⬭	⬭	⬭	⬭
Hands-on (non-programming or unplugged) activities	⬭	⬭	⬭	⬭	⬭
Programming activities	⬭	⬭	⬭	⬭	⬭
In-person teaching	⬭	⬭	⬭	⬭	⬭
Synchronous online teaching	⬭	⬭	⬭	⬭	⬭
Asynchronous online teaching	⬭	⬭	⬭	⬭	⬭

14. Any final comments?

Activity Feedback and Reflection

This form will be used to collect your reflections on activities as students progress through the activities. Please complete this form every day you use an activity from the How to Train Your Robot Curriculum.

1. Please enter your name or email.

   This is just to give me someway to track responses across multiple submissions.

   ____

2. Broadly, what topic did you cover with students?

   In the Educator Guide, these sections are labelled as 'Days'

   Check all that apply.

   - [ ] Section 1: Intro to AI and Ethics
   - [ ] Section 2: Intro to Algorithms, Image Recognition, and Algorithmic Bias
   - [ ] Section 3: Intro to Ethical Design and Text Classification
   - [ ] Section 4: AI Project Design
   - [ ] Section 5: Showcase Day!

3. Specifically, which activities did you complete?

Day 1

- Welcome
- What is AI?
- AI or Not?
- Ethical Dilemmas
- Intro to AI Blocks
- Reflect

Day 2

- Welcome
- Pizza Recipe Algorithm
- Pizza Delivery App
- Intro to Image Recognition
- Algorithmic Bias Discussion
- Animal Recognition
- Reflection

Day 3

- Welcome
- Ethical Matrices - Redesign Alexa
- Intro to Text Classification

Day 3 (cont.)

- Command Recognition
- Final Project Research
- Reflect

Day 4

- Welcome
- Final Project: Planning
- Final Project: Work Time
- Final Project: Peer Review
- Reflect

Day 5

- Welcome
- Final Project: Finishing Up
- Final Project: Showcase Preparation
- Final Project: Showcase
- Final Project: Reflect

4. Can you summarize the key skills and ideas that your students learned in today's lesson?

5. What was most engaging for your students?

6. What were some things that your students struggled with?

7. Did you adjust the activity by bringing something new to it or removing something?

8. Do you have any other comments about your experience conducting the activities?
   This is a place to put questions, advice for future teachers, ideas on how to improve the software/hardware design, or anything else that comes to mind.

Keep in touch!

If you have any questions or need help on the activities, feel free to reach out to Randi any time: randiw12@mit.edu.

E

AI for Wellbeing Teacher Surveys

Teacher Info Form

We are excited to have you join us this summer!

From July 24 - July 28, the MIT Personal Robots Group and MIT RAISE initiative are piloting a new AI literacy curriculum about AI for Wellbeing. This opportunity includes 5 hours of training plus 10 hours of co-teaching the curriculum to the S.T.E.A.M. Academy students over Zoom.

Please complete this form to **confirm** your participation in the program.

When: The professional development training will begin the weekend (exact date and time TBD) before July 24th. From July 24 - 27 teachers will need to be active from 9:30-12:30pm EST for teaching and training. The last session on July 28 will only be from 9:30-11:30am EST.

Stipend: We are offering a $500 stipend for teachers' participation in 5 hours of professional development plus 10 hours of co-teaching time.

Information about Research Study
This workshop is being run as a research study to contribute to the thesis of Randi Williams. Your participation in the research study portion of the work is *completely* voluntary and you maintain control of your data at all times. You can elect to withdraw your participation in the research study and to withdraw any data collected at any time, including after the workshop has been completed.

* Indicates required question

1. Email *

2. Please enter your first and last name. *

3. Please enter a mailing address where you can safely receive a package. *

All participants will be mailed a Jibo robot (www.youtube.com/watch?v=H0h20jRA5M0) for the duration of the workshop. We are collecting your mailing address so that we can send the robot to you.

Package details: medium sized, rectangular box that weighs about 6 pounds. Contains one (1) robot and packaging materials.

____

4. Please enter a phone number for delivery.

____

5. IGNORE THIS QUESTION, training will be on Sunday July 23 at 3pm.

When are you available for the first training session? [1 hour]

Please mark all times that you are available. We will get the session scheduled ASAP so that you don't have to hold the time on your calendar for too long.

Times are in EST.

Check all that apply.

	8-10am	10-12pm	12-2pm	2-4pm	4-6pm	6-8pm	Not available this day
Friday, July 21	☐	☐	☐	☐	☐	☐	☐
Saturday, July 22	☐	☐	☐	☐	☐	☐	☐
Sunday, July 23	☐	☐	☐	☐	☐	☐	☐

Research Consent Form

This section contains an informed consent form for participating in research. If you choose to have information collected for research, please fill in the two questions at the end of this section.

** You do not have to complete this form to participate in the workshop! **

Informed Consent Summary

The information below provides a summary of the research project you are being requested to participate in. Your participation in this research is voluntary and you can withdraw at any time.

- Purpose

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. You will guide children in interacting with a web-based platform and/or a physical robot during these educational interventions.

- Study Procedures

Please complete this survey about your background, perceptions of AI, and your thoughts about teaching AI in a classroom. As we go through the workshop, we will use anonymized forms to ask for your feedback on the content and approach to teaching AI. You may elect to participate in a follow-up interview after the workshop is over.

- Risks & Potential Discomfort

There are no anticipated risks greater than those associated with standard, virtual classroom environments.

You should read the information below, and ask questions about anything you do not understand before deciding whether or not to participate.

Purpose of This Study

The goal of this project is to evaluate education modules that teach children about different artificial intelligence concepts. Children will interact with a web-based platform and/or a physical robot during these educational interventions. Children’s understanding of artificial intelligence is presently under-researched. We hope to evaluate activities that help children understand how artificial intelligence works. In programing the robot, we hope that children will be able to gain an early understanding of the benefits and limitations of artificial intelligence. We hope that this will promote trust and provide a sound foundation for future technology education.

Study Procedures

If you volunteer to participate in this study, we would ask you to do the following things:

We will ask you questions about your background and any interactions with AI that you and your students have had as well as your perceptions of AI and teaching AI in middle school. Throughout the workshop, we will have you complete anonymous forms to give your feedback on modules. You may elect to participate in an interview following the workshop.

Risks and Benefits

Potential Risks and Discomforts

There are no anticipated risks greater than those associated with interaction with educational activities. If you no longer want to be a part of the study, then you can stop your survey or email us to delete your data.

Potential Benefits

Direct benefits to participants include a fun learning experience for you and hopefully your future students.

Payment for Participation

You will receive $500 for completing the professional development workshop. There is no further form of compensation given for participating in the study.

Privacy and Confidentiality

Any information obtained in connection with this study and that can be identified with you will remain confidential and will be disclosed only with your permission or as required by law. In addition, your information may be reviewed by authorized MIT representatives to ensure compliance with MIT policies and procedures.

Access to this data will be limited to the researchers who are performing the study. At any time, during or after, you can request that all data collected during your participation be destroyed.

Emergency Care and Compensation for Injury

If you feel that you have suffered an injury, which may include emotional trauma, as a result of participating in this study, please contact the person in charge of the study as soon as possible. In the event you suffer such an injury, M.I.T. may provide itself, or arrange for the provision of, emergency transport or medical treatment, including emergency treatment and follow-up care, as needed, or reimbursement for such medical services. M.I.T. does not provide any other form of compensation for injury. In any case, neither the offer to provide medical assistance, nor the actual provision of medical services shall be considered an admission of fault or acceptance of liability. Questions regarding this policy may be directed to MIT's Insurance Office, (617) 253-2823. Your insurance carrier may be billed for the cost of emergency transport or medical treatment if such services are determined not to be directly related to your participation in this study.

Identification of Investigators

If you have any questions or concerns about the research, please feel free to contact Randi Williams at randiw12@media.mit.edu (301) 633-8837 or Cynthia Breazeal at cynthiab@media.mit.edu.

6. Participation and Withdrawal

Your participation in this study is completely voluntary and you are free to choose whether you will be in it or not. If you choose to be in this study, you may subsequently withdraw at any time without penalty or consequences of any kind. The investigator may withdraw you from this research if circumstances arise.

Leave blank if not participating in research.

Mark only one oval.

- [ ] I understand that my participation is voluntary and that I may choose to withdraw at any time.

7. Rights of Research Subjects

You are not waiving any legal claims, rights or remedies because of your participation in this research study. If you feel you have been treated unfairly, or you have questions regarding your rights as a research subject, you may contact the Chairman of the Committee on the Use of Humans as Experimental Subjects, M.I.T., Room E25-143B, 77 Massachusetts Ave, Cambridge, MA 02139, phone 1-617-253 6787.

Leave blank if not participating in research.

Mark only one oval.

- I understand the procedures described above. My questions have been answered to my satisfaction, and I agree to participate in this study.

Pre-Questionnaire

We would like to start off by knowing how much you already know about computers and technology and why you are interested in teaching this topic.
If you don't know how to answer any of the questions, don't worry just do your best. We will not share your answers with others or tie answers back to your name or other identifying information

8. Have you heard of Artificial Intelligence (AI) before this opportunity?

Mark only one oval.

- Yes
- No
- Not sure

9. How would you define AI?

____

10. On a scale of 1 to 5, how much is technology a part of your classroom?

Mark only one oval.

	1	2	3	4	5	
I'm a	◯	◯	◯	◯	◯	I use technology often and actively seek new topics and tools. / not technology in the classroom

11. Have you ever taught AI to middle school students before?

Mark only one oval.

- ◯ Yes, I have taught a middle school AI course or workshop before.
- ◯ Yes, I have taught AI to middle school students as a one-off activity.
- ◯ No, I have taught AI before, but not to middle school students.
- ◯ No, I have never taught AI to students before.
- ◯ Other: ____________________

12. During this course students will become more familiar with AI concepts. What are some related skills and attitudes that you hope students develop through these activities?

____

____

____

____

____

13. Is there anything else that you would like to share about your experience as a teacher or motivation to teach AI?

Teacher Post-Questionnaire

Sign in to Google to save your progress. Learn more

To what extent do you agree with these statements?

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
I think AI is a subject that engages students with different interests.	○	○	○	○	○
I think it is important that students learn about AI at school.	○	○	○	○	○
My students were engaged in the AI workshop.	○	○	○	○	○
I consider myself to be very knowledgeable about AI.	○	○	○	○	○

Were any of your answers to the previous questions different before you participated in the AI workshop? If so, how did your answers change?

Your answer

Request edit access

Do you think that learning about AI should take time in the classroom? If so, what are skills students should develop? If not, why not?

Your answer

Summary of AI for Wellbeing

- Day 1: Students are exposed to the definition of AI, various examples, and potential benefits and harms. They begin programming the Jibo robot.

- Day 2: Students learn about Design Justice and how to use design to create more equitable spaces.

- Day 3:

Students learn about AI and non-AI text classification and the pros and cons of each.

- Day 4: Students explore AI for wellbeing and design their own interactive AI projects.

- Day 5: Students finish their projects and share their work.

What were the most important skills and ideas that your students learned in this workshop?

Your answer

What is one topic or skill that you would ADD to this curriculum?

Your answer

What is one topic or skill that you would REMOVE from this curriculum?

Your answer

Did this workshop change the way you would approach teaching Computer Science or AI to your students?

Your answer

Do you have any other feedback for us as we continue to work on this curriculum?

Your answer

Submit

Clear form

Daily Reflection

It's important to take time to reflect on the activities and students' learning and progress every day. This form will be used to collect your reflections.

Sign in to Google to save your progress. Learn more

Initials?

Your answer

Which day of the workshop is it?

Choose

Can you summarize the key skills and ideas that your students learned in today's lesson?

Your answer

What activities or moments were most engaging for students?

Your answer

What were some things that students struggled with?

Your answer

Request edit access

F

HTTYR v2.0 Student Surveys

Student Pre-Questionnaire - 2020-21

We would like to start off by knowing how much you already know about computers and technology and how interested you are in the topic.
If you don't know how to answer any of the questions, don't worry. Do your best. We will not share your answers with others.

Sign in to Google to save your progress. Learn more

What is your name?

Your answer

What is your teacher's name?

Your answer

Have you used block-based programming languages before? (like Scratch, Code.org, App Inventor)

	1	2	3	4	5	
I have never heard of these things before.	○	○	○	○	○	I am a master of block-based programming languages!

On a scale of 1 to 5, how much is technology a part of your life?

	1	2	3	4	5	
I'm afraid of technology	○	○	○	○	○	I can't live without technology

Request edit access

Have you heard of Artificial Intelligence (AI) before this workshop?

- ◯ Yes
- ◯ No
- ◯ Not sure

How would you describe Artificial Intelligence in your own words? Where have you seen the word used before?

Your answer

Which of the following statements is correct about machine learning (ML)?

- ◯ ML is a technique that enables computers to be trained with sample data to perform certain tasks.
- ◯ ML implies that humans will no longer have to think about how to solve problems, since computers will think for us (and better than us).
- ◯ ML implies that humans no longer need to learn to program.
- ◯ ML can only be implemented on supercomputers with enormous computing power.

Do you think that the following technologies use Artificial Intelligence?

	Yes	No	Not Sure
Google Search	○	○	○
Wireless Printer	○	○	○
Video Calls (FaceTime, Zoom)	○	○	○
YouTube "Up next" videos	○	○	○
Voice Assistants (Siri, Alexa, Google Home)	○	○	○
Snapchat Filters	○	○	○
GPS (Google Maps, Waze)	○	○	○
Washing Machines for clothes	○	○	○

Choose an example of technology (like one from above) that uses Artificial Intelligence and describe how it works.

Your answer

How much do the following statements describe you?

	Quite a lot	Somewhat	Not really	Not at all	Unsure
I don't know a lot about computers and technology.	○	○	○	○	○
I like to learn about computers and technology.	○	○	○	○	○
I usually find computers and technology confusing.	○	○	○	○	○
I would like a job with computers and technology.	○	○	○	○	○
Some of my friends or family work with computers and technology.	○	○	○	○	○
I think that AI is too hard to learn for most of the population.	○	○	○	○	○
I think that an evil artificial intelligence (AI) or robot will be developed, threatening the survival of humanity.	○	○	○	○	○

Student Post-Questionnaire - 2020-21

Thank you so much for being a part of this course. This is your final reflection and then you're 100% done!

If you don't know how to answer any of the questions, don't worry. Do your best. We will not share your answers with others.

Sign in to Google to save your progress. Learn more

What is your name?

Your answer

What is your teacher's name?

Your answer

How would you describe Artificial Intelligence in your own words?

Your answer

Request edit access

Which of the following statements is correct about machine learning (ML)?

- ○ ML is a technique that enables computers to be trained with sample data to perform certain tasks.
- ○ ML implies that humans will no longer have to think about how to solve problems, since computers will think for us (and better than us).
- ○ ML implies that humans no longer need to learn to program.
- ○ ML can only be implemented on supercomputers with enormous computing power.

Do you think that the following technologies use Artificial Intelligence?

	Yes	No	Not Sure
Google Search	○	○	○
Wireless Printer	○	○	○
Video Calls (FaceTime, Zoom)	○	○	○
YouTube "Up next" videos	○	○	○
Voice Assistants (Siri, Alexa, Google Home)	○	○	○
Snapchat Filters	○	○	○
GPS (Google Maps, Waze)	○	○	○
Washing Machines for clothes	○	○	○

Choose an example of technology that uses Artificial Intelligence and describe how it works.

Your answer

Request edit access

How much do the following statements describe you?

	Quite a lot	Somewhat	Not really	Not at all	Unsure
I don't know a lot about computers and technology.	○	○	○	○	○
I like to learn about computers and technology.	○	○	○	○	○
I usually find computers and technology confusing.	○	○	○	○	○
I would like a job with computers and technology.	○	○	○	○	○
Some of my friends or family work with computers and technology.	○	○	○	○	○
I think that AI is too hard to learn for most of the population.	○	○	○	○	○
I think that an evil artificial intelligence (AI) or robot will be developed, threatening the survival of humanity.	○	○	○	○	○

G

AI for Wellbeing Student Surveys

AI for Wellbeing Questionnaire

* Indicates required question

Enter your name. *

We will not share your answers with others.

Your answer

Part 1 - AI Knowledge

We would like to start off by exploring how much you already know about artificial intelligence. Don't worry if you don't know how to answer some of the questions. Just do your best.

Which of the following statements do you think is true about AI?

- ◯ AI is the simulation of different kinds of intelligence in computers and machines.
- ◯ Only machines that look like and act like humans are AIs.
- ◯ All robots are considered AI.
- ◯ The term "AI" describes systems, like Alexa or Siri, that can verbally interact with humans.

Which of the following statements is true about machine learning (ML)?

- ○ Without training data (an initial dataset collected by programmers) it is not possible to do machine learning.
- ○ With machine learning, computers learn to think the same way that a human being does.
- ○ The more data we use to train a machine learning system, the worse (more inaccurate) it becomes.
- ○ Machine learning does not need data to function, it learns automatically and does not depend on being fed information.

Which of the following are realistic benefits and harms of AI-powered chatbots

- ☐ Harm: Chatbots will create too many jobs and there won't be enough people to fill them.
- ☐ Benefit: Chatbots can talk with people, offering digital assistance whenever they need it.
- ☐ Benefit: Chatbots rarely make mistakes. They can work perfectly without humans checking them.
- ☐ Harm: A vehicle that is smart enough to have realistic conversations is likely smart enough to plan a revolt against humans.
- ☐ Harm: People might put too much trust in robots and use them in inappropriate ways.

Label the missing steps (A, B, and C) of AI-powered chatbots.

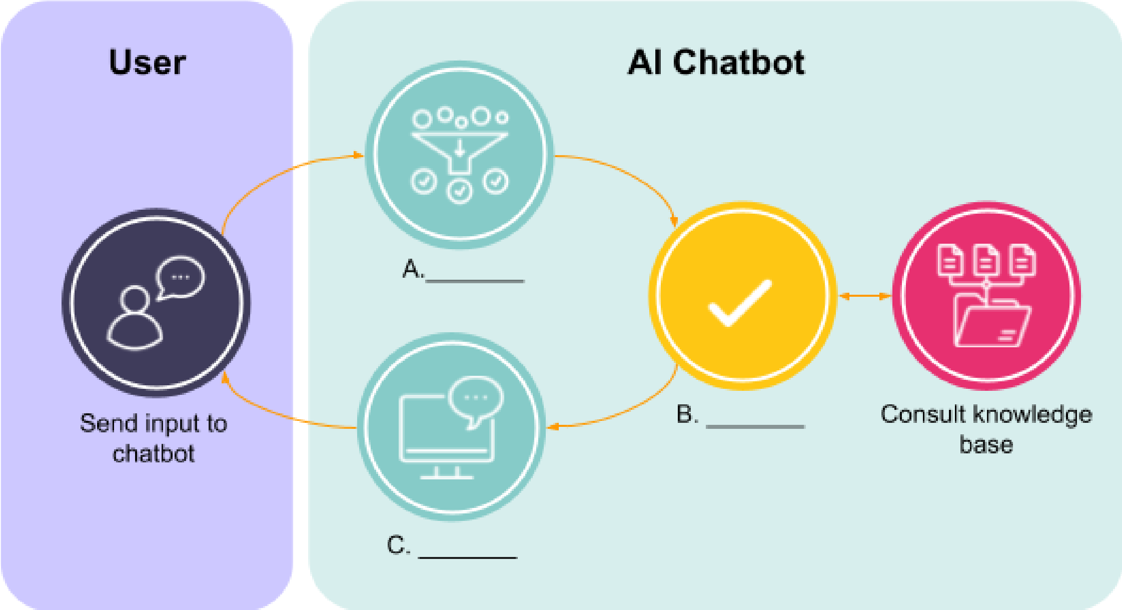

	A	B	C
Compose a reply	○	○	○
Analyze user input	○	○	○
Handle requests	○	○	○

Request edit access

Which of these tasks do chatbots do using AI?

	AI	Not AI	Unsure
Analyze the emotional tone in a user's input.	◯	◯	◯
Recognize what task a user wants by comparing their input to examples.	◯	◯	◯
Send requests to the Internet to get information.	◯	◯	◯
Use a microphone to capture what a user says.	◯	◯	◯
Learn how to respond appropriately after getting feedback from users.	◯	◯	◯

Which of the following statements are true about text classifiers?

	True	False	Unsure
It is easy to understand how neural network classifiers make decisions.	○	○	○
It is easy to understand how rule-based text classifiers make decisions.	○	○	○
People can usually build neural network classifiers faster and more easily than rule-based classifiers.	○	○	○
Rule-based text classifiers are good at dealing with noisy and unclear input.	○	○	○
You need a lot of expert knowledge to create rule-based classifiers.	○	○	○

Word vectors are arrays of numbers that represent the meaning of a word. They make it possible to do mathematical operations on words, e.g. king - man + woman = queen. Have you heard of this concept before?

- ◯ I've never heard of word vectors.
- ◯ I've heard of word vectors but don't really understand them.
- ◯ Yes, and I understand how word vectors work.

How do you think word vectors created?

- ◯ Word vectors are calculated by looking at the letters in words.
- ◯ Word vectors are created by assigning a number to every word in the dictionary.
- ◯ Every word already has a word vector associated with it that can be readily accessed.
- ◯ Word vectors are made by comparing many examples of what words mean and how they are used.

How might word vectors be impacted by bias?

- ◯ Word vectors are trained on large datasets where they may learn things that humans would consider biased.
- ◯ Word vectors are directly trained by people and can learn bias from their biases.
- ◯ Language and words contain biases that are then reflected in word vectors.
- ◯ Choosing the wrong equations when calculating word vector can create bias.

Request edit access

How familiar are you with the following skills?

	1 - I've never heard of this	2 - I've heard of this, but don't know how to do it	3 - I understand this well and can do it on my own
Programming (Scratch, Code.org, Python) and computational thinking (loops, if...then, abstraction)	○	○	○
Using AI to have a positive societal impact	○	○	○
Weighing the consequences of AI systems	○	○	○
Presenting your AI knowledge to an audience	○	○	○

Part 2 - Perspectives about AI

Now, we want to know what you think about AI and how interested you are in the topic. There are no wrong answers, just be honest.

Which of the following words would you use to describe AI? Select up to 4.

- [ ] Intelligent
- [ ] Unfair
- [ ] Impressive
- [ ] Error-prone
- [ ] Manipulative
- [ ] Dangerous
- [ ] Creative
- [ ] Useful
- [ ] Beneficial
- [ ] Snooping

How much do you agree with the following statements?

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Creating and thinking about AI is only for people who are good with computers.					
I can use AI applications to make my everyday life easier.					
I can keep up with the latest innovations in AI applications.					
I see people I can relate to being successful in the field of AI.					
I can rely on my skills to solve difficult AI problems.					

Request edit access

How much do you agree with the following statements?

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I'd like to learn AI so that I can transform society for the better.	○	○	○	○	○
I believe it is important to talk to members of my household and community about AI	○	○	○	○	○
When learning about new AI technology I weigh the consequences of that technology.	○	○	○	○	○
I feel anxious while doing activities about computers and AI.	○	○	○	○	○
I would be willing to learn more about AI because it has some value to me.	○	○	○	○	○
I'd like to use my AI knowledge to help my community.	○	○	○	○	

Request edit access

Do you see yourself as someone who can be successful in AI? Why or why not?

Your answer

Submit Page 1 of 1 Clear form

AI for Wellbeing Questionnaire

* Indicates required question

1. Enter your name. *

   We will not share your answers with others. (Read aloud)

   ____

Part 1 - AI Knowledge

We would like to start off by exploring how much you already know about artificial intelligence. Don't worry if you don't know how to answer some of the questions. Just do your best. (Read aloud)

2. Which of the following statements do you think is true about AI? (Read aloud)

   Mark only one oval.

   - ◯ A. AI is the simulation of different kinds of intelligence in computers and machines.
   - ◯ B. All robots are considered AI.
   - ◯ C. The term "AI" describes systems, like Alexa or Siri, that can verbally interact with humans.
   - ◯ D. Only machines that look like and act like humans are AIs.

3. Select whether you think the following examples uses AI or not.

Mark only one oval per row.

	AI	Not AI	Unsure
Virtual assistants	⬭	⬭	⬭
Phone teller menus	⬭	⬭	⬭
Non-player characters (NPCs) in video games	⬭	⬭	⬭

4. Which of the following statements is true about machine learning (ML)? (Read aloud)

Mark only one oval.

- ⬭ A. Without training data (an initial dataset collected by programmers) it is not possible to do machine learning.
- ⬭ B. With machine learning, computers learn to think the same way that a human being does.
- ⬭ C. The more data we use to train a machine learning system, the worse (more inaccurate) it becomes.
- ⬭ D. Machine learning does not need data, it learns automatically and does not depend on being fed information.

5. Which of the following are realistic benefits and harms of AI-powered chatbots? (Read aloud)

Check all that apply.

☐ A. Harm: Chatbots will create too many jobs and there won't be enough people to fill them.

☐ B. Benefit: Chatbots can talk with people, offering digital assistance whenever they need it.

☐ C. Benefit: Chatbots rarely make mistakes. They can work perfectly without humans checking them.

☐ D. Harm: A vehicle that is smart enough to have realistic conversations is likely smart enough to plan a revolt against humans.

☐ E. Harm: People might put too much trust in robots and use them in inappropriate ways.

6. Label the missing steps (A, B, and C) of AI-powered chatbots. (Read aloud)

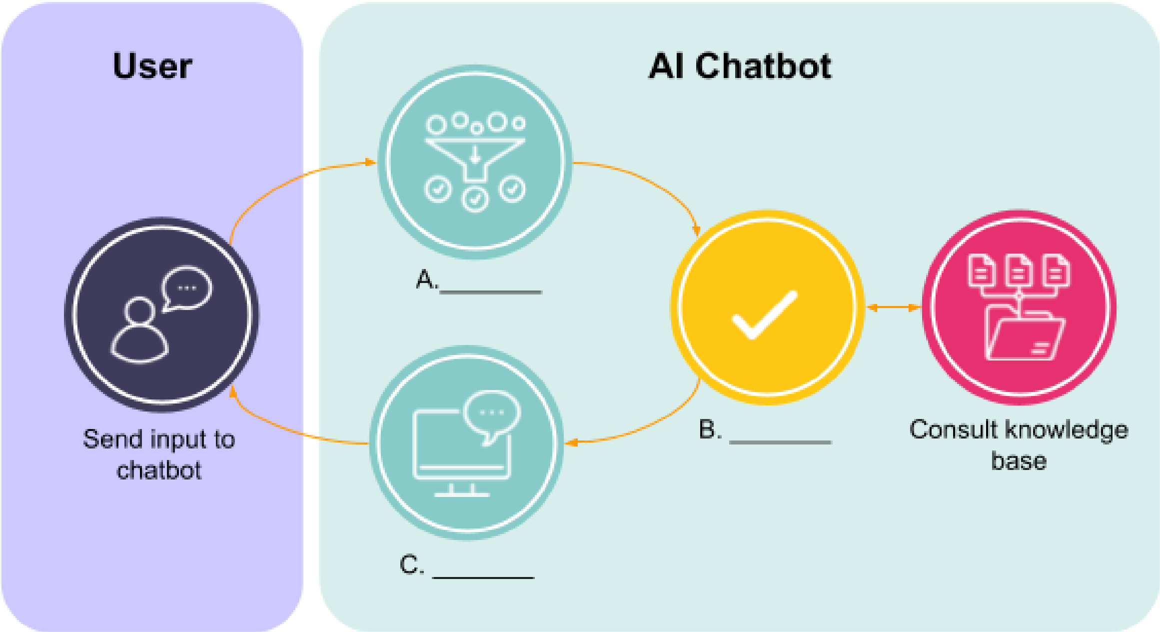

Mark only one oval per row.

	A	B	C
Write a reply	◯	◯	◯
Understand the user's message	◯	◯	◯
Handle requests	◯	◯	◯

7. How familiar are you with the following skills on a scale of 1 (I've never heard of this) to 3 (I understand this well). (Read aloud)

Mark only one oval per row.

	1 - I've never heard of this	2 - I've heard of this, but don't know how to do it	3 - I understand this well and can do it on my own
Programming (Scratch, Code.org, Python) and computational thinking (loops, if...then, abstraction)	◯	◯	◯
Using AI to have a positive societal impact	◯	◯	◯
Weighing the consequences of AI systems	◯	◯	◯
Presenting your AI knowledge to an audience	◯	◯	◯

Part 2 - Perspectives about AI (Read aloud)

Now, we want to know what you think about AI and how interested you are in the topic. There are no wrong answers, just be honest.

8. Which of the following words would you use to describe AI? Select up to 4. (Read aloud)

Check all that apply.

- [ ] Intelligent
- [ ] Unfair
- [ ] Impressive
- [ ] Error-prone
- [ ] Manipulative
- [ ] Dangerous
- [ ] Creative
- [ ] Useful
- [ ] Beneficial
- [ ] Snooping

9. How much do you agree with the following statements? (Read aloud)

Mark only one oval per row.

	1 - Strongly disagree	2 - Disagree	3 - Neutral	4 - Agree	5 - Strongly agree
I can use AI applications to make my everyday life easier.	⬭	⬭	⬭	⬭	⬭
I can rely on my skills to solve difficult AI problems.	⬭	⬭	⬭	⬭	⬭
I can keep up with the latest innovations in AI applications.	⬭	⬭	⬭	⬭	⬭
Creating and thinking about AI is only for people who are good with computers.	⬭	⬭	⬭	⬭	⬭
I see people I can relate to being successful in the field of AI.	⬭	⬭	⬭	⬭	⬭

10. How much do you agree with the following statements? (Read aloud)

Mark only one oval per row.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
When learning about new AI technology I weigh the consequences of that technology.	⬭	⬭	⬭	⬭	⬭
I feel anxious while doing activities about computers and AI.	⬭	⬭	⬭	⬭	⬭
I believe it is important to talk to members of my household and community about AI	⬭	⬭	⬭	⬭	⬭
I would be willing to learn more about AI because it has some value to me.	⬭	⬭	⬭	⬭	⬭
I'd like to use my AI knowledge to help my community.	⬭	⬭	⬭	⬭	⬭
I'd like to learn AI so that I					

I'd like to learn AI so that I can transform society for the better.					

11. Did this workshop make you feel more confident about being successful in AI? If yes, how? If not, why not? ([Read aloud](#))

12. Upload your .sb3 project files from your final project or anything else you built during the week.

If you have more than five items to upload, e-mail them to randiw12@mit.edu.

Files submitted:

H

Sparki Pilot Surveys

LLMs for Education Workshop

* Indicates required question

Pre-Assessment

1. Please enter the username given with your computer

2. How familiar are you with the following terms?

Mark only one oval per row.

	1 - I've never heard of this	2 - I've heard of this but only understand it a little	3 - I am beginning to understand this, but still need some help	4 - I understand this well and can talk about it without help	5 - I understand this very well and can teach it to someone else
Artificial Intelligence					
Chatbots					
Large language model					

3. How much do you agree or disagree with the following statements about your ability to develop creative ideas?

Mark only one oval per row.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I trust my creative abilities	⬭	⬭	⬭	⬭	⬭
I am good at coming up with new ideas	⬭	⬭	⬭	⬭	⬭
I have a good imagination	⬭	⬭	⬭	⬭	⬭
I am sure I can deal with problems requiring creative thinking	⬭	⬭	⬭	⬭	⬭

4. How much do you agree or disagree with the following statements about using AI?

Mark only one oval per row.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I can use AI meaningfully to achieve my everyday goals	⬭	⬭	⬭	⬭	⬭
I can work together with AI in everyday life	⬭	⬭	⬭	⬭	⬭
I know the most important concepts of AI	⬭	⬭	⬭	⬭	⬭
I know the defintion of AI	⬭	⬭	⬭	⬭	⬭
I can assess what the limitations and opportunities of using AI are	⬭	⬭	⬭	⬭	⬭
I can think of new uses and imagine possible future uses of AI	⬭	⬭	⬭	⬭	⬭

5. What grade in school are you in (for the upcoming 2023-24 school year)?

____

6. What is your gender?

Mark only one oval.

- Female
- Male
- Non-binary/Genderqueer
- Prefer not to say
- Other: _______________________________________

Pause

Once you get to this point, pause until we continue to the next part of the workshop.

Post-Assessment

7. Did you copy and paste your slides into the shared Google Presentation? <Insert link> *

Mark only one oval.

- Yes

8. How much do you agree or disagree with the following statements about using AI?

Mark only one oval per row.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I can use AI meaningfully to achieve my everyday goals	⬭	⬭	⬭	⬭	⬭
I can work together with AI in everyday life	⬭	⬭	⬭	⬭	⬭
I know the most important concepts of AI	⬭	⬭	⬭	⬭	⬭
I know the defintion of AI	⬭	⬭	⬭	⬭	⬭
I can assess what the limitations and opportunities of using AI are	⬭	⬭	⬭	⬭	⬭
I can think of new uses and imagine possible future uses of AI	⬭	⬭	⬭	⬭	⬭

9. In what ways was Sparki helpful?

Check all that apply.

- [ ] Coming up with new ideas for my project
- [ ] Correcting my grammar or spelling
- [ ] Connecting my ideas with different fields or subjects
- [ ] Exploring alternative ideas I hadn't thought about
- [ ] Overcoming creative blocks
- [ ] Other: ____________________

10. Do you have any other feedback about ways that Sparki was helpful for you?

____

____

____

____

____

11. How much do you agree with the following statements about the usability of Sparki

Mark only one oval per row.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Sparki was exciting or engaging	⬭	⬭	⬭	⬭	⬭
Sparki was difficult to use	⬭	⬭	⬭	⬭	⬭
Someone can learn how to use Sparki quickly	⬭	⬭	⬭	⬭	⬭
I feel confident showing someone else how to use Sparki	⬭	⬭	⬭	⬭	⬭
Using Sparki made my task easier	⬭	⬭	⬭	⬭	⬭

12. Do you have any other feedback about how Sparki could be improved?

I

Sparki Experiment Surveys

Consent Form

Hello, welcome to the **Sparki User Study**!

In this study, you will use an AI tool called Sparki to come up with AI project ideas. This study is part of a university research project. Your decision to complete this study is voluntary. We will record your responses to this survey and log your interactions with the Sparki tool. The only other information we will have is when you completed the survey. There is no way for us to identify you. The research results may be presented at scientific meetings or published in scientific journals.

If you agree to participate in this study, please select 'I agree' to continue. Otherwise select 'I do not agree' to return your study to Prolific.

- ○ I agree to participate in this study.
- ○ I do not agree to participate in this study.

Introduction

Please verify that your Participant ID from Prolific is properly copied in the textbox below.

${e://Field/username}

Participant Background

What is your gender identity?

- ☐ Non-binary / third gender
- ☐ Genderqueer
- ☐ Woman
- ☐ Man
- ☐ ________ A gender identity not listed here (please specify)
- ☐ Prefer not to say

Please select your age group

- ○ 18 to 24
- ○ 25 to 29
- ○ 30 to 39
- ○ 40 to 49
- ○ 50 to 59
- ○ 60 to 69
- ○ 70 plus

Which of the following best describes your fluency with the English language?

- ○ English is my first language
- ○ I am fluent in English
- ○ My English is intermediate or conversational
- ○ I use limited English

On a scale of 1 to 5, please rate your preferences for working in the following conditions.

	Not at all	Somewhat	Moderately	Quite a bit	Very much
I prefer working alone, without direct interaction from others.	○	○	○	○	○
I prefer collaborating and sharing ideas with others.	○	○	○	○	○

Familiarity with Topic

How familiar are you with Artificial Intelligence (AI)

- ○ Extremely familiar (expert)
- ○ Very familiar
- ○ Moderately familiar
- ○ Somewhat familiar
- ○ Not at all familiar

Have you ever used an AI powered technology (e.g., ChatGPT, Dall-E, Stable Diffusion, GPT3, Bing AI, etc.) before?

- ○ Yes
- ○ No
- ○ I'm not sure

If you answered yes to the previous question, please describe how you have used AI powered technologies in the past. (Leave blank if you are not sure.)

Expectations Pre-test Questions

On a scale of 1 to 5, please rate how helpful you think an AI tool could be in helping you brainstorm and write about a project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
An AI tool can help me write better.	○	○	○	○	○
An AI tool can help me write more quickly.	○	○	○	○	○
An AI tool can help me clarify and expand on my ideas.	○	○	○	○	○
An AI tool can help me explore new ideas.	○	○	○	○	○

On a scale of 1 to 5, please rate the ways in which you think an AI tool could help you brainstorm and write about a project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
An AI tool could provide helpful suggestions.	○	○	○	○	○
An AI tool could provide relevant suggestions.	○	○	○	○	○
An AI tool could provide creative suggestions.	○	○	○	○	○

Interaction with Sparki - Project 1, non-AI

Now it is time for you to brainstorm a novel AI-powered project idea. You can pick any topic, but remember not to include personal information about yourself or others in the project.

Sparki Tutorial Video

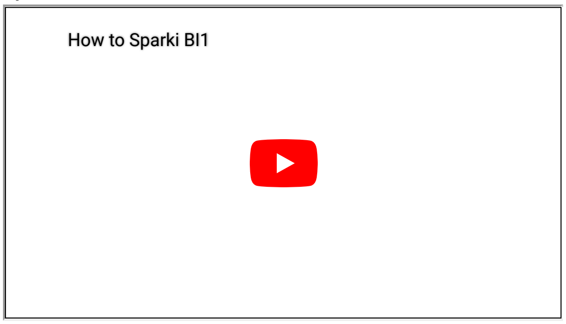

To open Sparki, follow the link below and enter the following credentials:
Key: sparki-${e://Field/username}

[Click Here to Launch SPARKI]

Please note that your work in Sparki will be recorded. Interaction with the agent will be recorded (text transcripts). Failure to write and submit a completed project (with all fields filled out) will result in invalidation of the task and no payment.

Double check! Did you click the link above and finish writing your project idea in Sparki notes?

- ○ No
- ○ Yes

Expectations Post-test Questions, non-AI

On a scale of 1 to 5, rate how helpful Sparki notes were in helping you brainstorm and write about your project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
Sparki notes helped me write better.	○	○	○	○	○
Sparki notes helped me write more quickly.	○	○	○	○	○
Sparki notes helped me clarify and expand on my ideas.	○	○	○	○	○
Sparki notes helped me explore new ideas.	○	○	○	○	○

On a scale of 1 to 5, please rate the ways in which Sparki notes helped you brainstorm and write about a project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
Sparki notes provided helpful suggestions.	○	○	○	○	○
Sparki notes provided relevant suggestions.	○	○	○	○	○
Sparki notes provided creative suggestions.	○	○	○	○	○

Which of the following statements are true about how you used Sparki notes? (You may pick more than one).

- ☐ I was collaborating with Sparki notes.
- ☐ I was just copying what Sparki notes told me to write.
- ☐ I looked at the information in Sparki notes but did not use it.

Please elaborate on what you enjoyed about using Sparki notes, if anything.

Please elaborate on what downsides there were to using Sparki notes, if any.

Interaction with Sparki - Project 2, AI

Now it is time for you to brainstorm a novel AI-powered project idea. You can pick any topic, but remember not to include personal information about yourself or others in the project.

Sparki Tutorial Video

To open Sparki, follow the link below and enter the following credentials:
Key: sparki-${e://Field/username}

Click Here to Launch SPARKI

Please note that your work in Sparki will be recorded. Interaction with the agent will be recorded (text transcripts). Failure to write and submit a completed project (with all fields filled out) will result in invalidation of the task and no payment.

Double check! Did you click the link above and finish writing your project idea with Sparki?

- ○ No
- ○ Yes

Expectations Post-test Questions, AI

On a scale of 1 to 5, rate how helpful Sparki was in helping you brainstorm and write about your project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
Sparki helped me write better.	○	○	○	○	○
Sparki helped me write more quickly.	○	○	○	○	○
Sparki helped me clarify and expand on my ideas.	○	○	○	○	○
Sparki helped me explore new ideas.	○	○	○	○	○

On a scale of 1 to 5,Rate your level of agreement with how Sparki helped you brainstorm and write your project idea.

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
Sparki provided helpful suggestions.	○	○	○	○	○
Sparki provided relevant suggestions.	○	○	○	○	○
Sparki provided creative suggestions.	○	○	○	○	○

Which of the following statements are true about how you used Sparki? (You may pick more than one).

- ☐ I was collaborating with Sparki.
- ☐ I was just copying what Sparki told me to write.
- ☐ I looked at Sparki's suggestions but did not use them.

Please elaborate on what you enjoyed about using Sparki, if anything.

Please elaborate on what downsides there were to using Sparki, if any.

Interaction with Sparki - Project 1, AI

Now it is time for you to brainstorm a novel AI-powered project idea. You can pick any topic, but remember not to include personal information about yourself or others in the project.

Sparki Tutorial Video

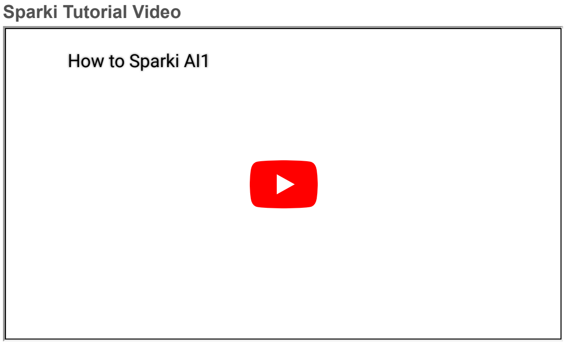

To open Sparki, follow the link below and enter the following credentials:
Key: sparki-${e://Field/username}

[Click Here to Launch SPARKI]

Please note that your work in Sparki will be recorded. Interaction with the agent will be recorded (text transcripts). Failure to write and submit a completed project (with all fields filled out) will result in invalidation of the task and no payment.

Double check! Did you click the link above and finish writing your project idea with Sparki?

○ No
○ Yes

Interaction with Sparki - Project 2, non-AI

Now it is time for you to brainstorm a novel AI-powered project idea. You can pick any topic, but remember not to include personal information about yourself or others in the project.

Sparki Tutorial Video

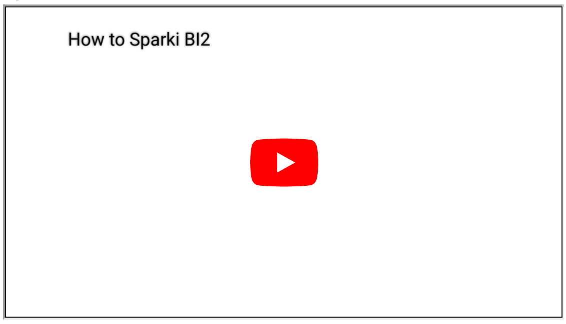

To open Sparki, follow the link below and enter the following credentials:
Key: sparki-${e://Field/username}

[Click Here to Launch SPARKI]

Please note that your work in Sparki will be recorded. Interaction with the agent will be recorded (text transcripts). Failure to write and submit a completed project (with all fields filled out) will result in invalidation of the task and no payment.

Double check! Did you click the link above and finish writing your project idea in Sparki notes?

○ No
○ Yes

Whole Study Post-Test Questions

On a scale of 1 to 5, please rate your experience with AI:

	Strongly disagree	Disagree	Neutral / Neither agree nor disagree	Agree	Strongly agree
I can work together and communicate with AI to achieve my everyday goals	○	○	○	○	○
I can think of new uses and imagine possible future uses of AI	○	○	○	○	○
I have a better understanding of AI	○	○	○	○	○

On a scale of 1 to 5, rate where you fall in comparing Sparki notes (information cards) to the Sparki AI tool (chatbot).

Compared to Sparki notes, ...

	Much worse	Worse	About the same	Better	Much better
...the Sparki AI tool made the quality of my writing	○	○	○	○	○

	Much worse	Worse	About the same	Better	Much better
...the Sparki AI tool made the expression of my ideas	○	○	○	○	○
...the Sparki AI tool made the swiftness of my writing	○	○	○	○	○
...the Sparki AI tool made my exploration of new ideas	○	○	○	○	○

If you did this project again, would you prefer to work with Sparki or Sparki notes? Please elaborate.

Do you have any other thoughts about this study? Feel free to elaborate on any of your previous answers here.

End of Survey message

Thank you for taking part in this study. Please click the button below to be redirected back to Prolific and register your submission.

Bibliography

Here are the references in citation order.

[1] Insider Intelligence. 'Voice Assistants in 2023: Usage, growth, and future of the AI voice assistant market'. In: (2023) (cited on page 12).

[2] Similarweb. *chat.openai.com*. Accessed: 2023-11-05. 2023. URL: https://www.similarweb.com/website/chat.openai.com/#overview (cited on page 12).

[3] Edith K Ackermann. 'Playthings that do things: a young kid's" incredibles"!' In: *Proceedings of the 2005 conference on Interaction design and children*. 2005, pp. 1–8 (cited on page 13).

[4] Peter H Kahn Jr et al. '"Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot.' In: *Developmental psychology* 48.2 (2012), p. 303 (cited on page 13).

[5] Sherry Turkle. 'Technology and human vulnerability. A conversation with MIT's Sherry Turkle.' In: *Harvard Business Review* 81.9 (2003), pp. 43–50 (cited on page 13).

[6] Sherry Turkle et al. 'Encounters with kismet and cog: Children respond to relational artifacts'. In: *Digital media: Transformations in human communication* 120 (2006) (cited on page 13).

[7] Gallup and Northeastern University. 'Facing the future: US, UK, and Canadian citizens call for a unified strategy for the AI age'. In: (2019) (cited on page 13).

[8] Stefania Druga et al. '"Hey Google is it ok if I eat you?" Initial explorations in child-agent interaction'. In: *Proceedings of the 2017 conference on interaction design and children*. ACM. 2017, pp. 595–600 (cited on page 13).

[9] Emily McReynolds et al. 'Toys that Listen: A Study of Parents, Children, and Internet-Connected Toys'. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM. 2017, pp. 5197–5207 (cited on page 13).

[10] V. Rideout and M. B. Robb. 'Social media social life'. In: (2018). [Online; accessed May 2022] (cited on page 13).

[11] Randi Williams. 'PopBots: leveraging social robots to aid preschool children's artificial intelligence education'. MA thesis. Massachusetts Institute of Technology, 2018 (cited on pages 13, 125).

[12] Code.org, CSTA, and ECEP Alliance. '2022 State of Computer Science Education: Understanding Our National Imperative'. In: (2022) (cited on pages 14, 23, 127).

[13] Dario Cvencek, Andrew N Meltzoff, and Anthony G Greenwald. 'Math–gender stereotypes in elementary school children'. In: *Child development* 82.3 (2011), pp. 766–779 (cited on page 14).

[14] Amanda Sullivan and Marina Umashi Bers. 'Girls, boys, and bots: Gender differences in young children's performance on robotics and programming tasks'. In: *Journal of Information Technology Education. Innovations in Practice* 15 (2016), p. 145 (cited on pages 14, 101).

[15] Yolanda A Rankin, Jakita O Thomas, and Sheena Erete. 'Real talk: Saturated sites of violence in CS education'. In: *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 2021, pp. 802–808 (cited on page 14).

[16] Kapor Center. 'Culturally responsive-sustaining computer science education: A framework'. In: 2021 (cited on pages 14, 23, 24).

[17] Alexandrea R Golden et al. 'What was a gap is now a chasm: Remote schooling, the digital divide, and educational inequities resulting from the COVID-19 pandemic'. In: *Current Opinion in Psychology* (2023), p. 101632 (cited on page 14).

[18] Ruha Benjamin. *Race after technology: Abolitionist tools for the new jim code*. 2019 (cited on pages 15, 23).

[19] Safiya Umoja Noble. 'Algorithms of oppression'. In: *Algorithms of Oppression*. New York University Press, 2018 (cited on pages 15, 23).

[20] Sarah Myers West, Meredith Whittaker, and Kate Crawford. 'Discriminating systems'. In: *AI Now* (2019) (cited on pages 15, 23).

[21] Tolga Bolukbasi et al. 'Man is to computer programmer as woman is to homemaker? debiasing word embeddings'. In: *Advances in neural information processing systems* 29 (2016), pp. 4349–4357 (cited on page 15).

[22] Virginia Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press, 2018 (cited on page 15).

[23] Cathy O'Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway books, 2016 (cited on page 15).

[24] Katlyn Turner, Danielle Wood, and Catherine D'Ignazio. 'The abuse and misogynoir playbook'. In: *The State of AI Ethics Report January* (2021), pp. 15–34 (cited on page 16).

[25] Harini Suresh and John V Guttag. 'A framework for understanding unintended consequences of machine learning'. In: *arXiv preprint arXiv:1901.10002* 2.8 (2019) (cited on pages 16, 19, 30, 42).

[26] Kimberly A Scott, Kimberly M Sheridan, and Kevin Clark. 'Culturally responsive computing: A theory revisited'. In: *Learning, Media and Technology* 40.4 (2015), pp. 412–436 (cited on pages 16, 23, 24, 43).

[27] CAST. 'Universal Design for Learning Guidelines version 2.2'. In: (2018) (cited on pages 16, 21, 22, 56).

[28] Sasha Costanza-Chock. *Design justice: Community-led practices to build the worlds we need*. The MIT Press, 2020 (cited on pages 16, 67, 105).

[29] Blakeley H Payne. 'Can my algorithm be my opinion?: an AI+ ethics curriculum for middle school students'. MA thesis. Massachusetts Institute of Technology, 2020 (cited on pages 17, 22, 25, 36, 43, 56, 59).

[30] Joy Buolamwini. 'Facing the Coded Gaze with Evocative Audits and Algorithmic Audits'. PhD thesis. Massachusetts Institute of Technology, 2022 (cited on pages 17, 43).

[31] Jean J Ryoo, Jane Margolis, and Allison Scott. 'Begin Again: Why CS Education Must be Reimagined'. In: *2021 Conference on Research in Equitable and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*. IEEE. 2021, pp. 1–2 (cited on pages 17, 19, 25, 43).

[32] Beverly Daniel Tatum. 'Why are all the Black kids sitting together in the cafeteria'. In: *And Other Conversations About Race, New York, Basic* (1997) (cited on page 17).

[33] David Touretzky et al. 'Envisioning AI for K-12: What should every child know about AI?' In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 33. 01. 2019, pp. 9795–9799 (cited on pages 18, 25, 33, 34, 39, 42, 56, 125).

[34] Duri Long and Brian Magerko. 'What is AI literacy? Competencies and design considerations'. In: *Proceedings of the 2020 CHI conference on human factors in computing systems*. 2020, pp. 1–16 (cited on pages 18, 23, 35, 42, 56, 125).

[35] UNESCO. 'Draft AI competency frameworks for teachers and for school students'. In: 2023 (cited on pages 18, 56).

[36] Karen Brennan and Mitchel Resnick. 'New frameworks for studying and assessing the development of computational thinking'. In: *Proceedings of the 2012 annual meeting of the American educational research association, Vancouver, Canada*. Vol. 1. 2012, p. 25 (cited on pages 19, 31, 33, 42, 44, 48).

[37] Natalie Lao. 'Reorienting Machine Learning Education Towards Tinkerers and ML-Engaged Citizens'. PhD thesis. Massachusetts Institute of Technology Cambridge, MA, USA, 2020 (cited on pages 19, 31, 49, 125).

[38] Harini Suresh and Natalie Lao. 'Opening the Black Box: Educational Machine Learning Videos for a General Public Audience'. In: (2020) (cited on page 19).

[39] Tia C Madkins et al. 'Culturally relevant computer science pedagogy: From theory to practice'. In: *2019 research on equity and sustained participation in engineering, computing, and technology (RESPECT)*. IEEE. 2019, pp. 1–4 (cited on pages 20, 23, 24).

[40] Davy Tsz Kit Ng et al. 'Conceptualizing AI literacy: An exploratory review'. In: *Computers and Education: Artificial Intelligence* 2 (2021), p. 100041 (cited on pages 20, 21).

[41] Maria-Jose Ojeda-Bazaran et al. 'State of the Art of Teaching-Learning of Artificial Intelligence at Early Ages'. In: *2021 XVI Latin American Conference on Learning Technologies (LACLO)*. IEEE. 2021, pp. 165–172 (cited on page 20).

[42] Jiahong Su and Weipeng Yang. 'Artificial intelligence in early childhood education: A scoping review'. In: *Computers and Education: Artificial Intelligence* (2022), p. 100049 (cited on page 20).

[43] Saman Rizvi, Jane Waite, and Sue Sentance. 'Artificial Intelligence teaching and learning in K-12 from 2019 to 2022: A systematic literature review'. In: *Computers and Education: Artificial Intelligence* (2023), p. 100145 (cited on page 20).

[44] Ismaila Temitayo Sanusi et al. 'Survey of resources for introducing machine learning in K-12 context'. In: *2021 IEEE Frontiers in Education Conference (FIE)*. IEEE. 2021, pp. 1–9 (cited on pages 20, 22, 31, 102).

[45] Christiane Gresse von Wangenheim et al. 'Visual tools for teaching machine learning in K-12: A ten-year systematic mapping'. In: *Education and Information Technologies* 26.5 (2021), pp. 5733–5778 (cited on page 20).

[46] Irene Lee et al. 'Computational thinking for youth in practice'. In: *Acm Inroads* 2.1 (2011), pp. 32–37 (cited on page 20).

[47] Chris-Bennet Fleger, Yousuf Amanuel, and Johannes Krugel. 'Learning Tools Using Block-based Programming for AI Education'. In: *2023 IEEE Global Engineering Education Conference (EDUCON)*. IEEE. 2023, pp. 1–5 (cited on pages 20, 22).

[48] Seymour A Papert. *Mindstorms: Children, computers, and powerful ideas*. Basic books, 2020 (cited on pages 21, 26, 63).

[49] Mitchel Resnick. *Lifelong kindergarten: Cultivating creativity through projects, passion, peers, and play*. MIT press, 2017 (cited on page 21).

[50] Mike Tissenbaum, Josh Sheldon, and Hal Abelson. 'From computational thinking to computational action'. In: *Communications of the ACM* 62.3 (2019), pp. 34–36 (cited on page 21).

[51] Joel Michael and Harold I Modell. *Active learning in secondary and college science classrooms: A working model for helping the learner to learn*. Routledge, 2003 (cited on pages 21, 22, 55).

[52] Charles C Bonwell and James A Eison. *Active learning: Creating excitement in the classroom. 1991 ASHE-ERIC higher education reports.* ERIC, 1991 (cited on pages 21, 22, 55).

[53] Laura Kieran and Christine Anderson. 'Connecting universal design for learning with culturally responsive teaching'. In: *Education and Urban Society* 51.9 (2019), pp. 1202–1216 (cited on page 21).

[54] David H Rose and Anne Meyer. *A practical reader in universal design for learning.* ERIC, 2006 (cited on pages 21, 56).

[55] Ken Kahn et al. 'Deep learning programming by all'. In: (2020) (cited on pages 21, 35).

[56] Hannah H. Pang. 'Computational Action in Action: Process and Tools that Empower Students to Make a Real-world Impact Using Technology'. MA thesis. Massachusetts Institute of Technology, 2022 (cited on page 21).

[57] Jessica Van Brummelen. 'Empowering K-12 Students to Understand and Design Conversational Agents: Concepts, Recommendations and Development Platforms'. PhD thesis. Massachusetts Institute of Technology, 2022 (cited on page 21).

[58] Safinah Ali, Daniella DiPaola, and Cynthia Breazeal. 'What are GANs?: Introducing Generative Adversarial Networks to Middle School Students'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 17. 2021, pp. 15472–15479 (cited on pages 22, 35, 40).

[59] Safinah Ali et al. 'Exploring Generative Models with Middle School Students'. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–13 (cited on pages 22, 35, 36, 40, 49).

[60] Irene Lee et al. 'Developing Middle School Students' AI Literacy'. In: *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 2021, pp. 191–197 (cited on pages 22, 35, 36, 40, 45, 48, 49).

[61] Ziv Epstein et al. 'Interpolating GANs to scaffold autotelic creativity'. In: *arXiv preprint arXiv:2007.11119* (2020) (cited on page 22).

[62] David Bau et al. 'Seeing what a gan cannot generate'. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, pp. 4502–4511 (cited on page 22).

[63] Phillip Isola et al. 'Image-to-image translation with conditional adversarial networks'. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 1125–1134 (cited on page 22).

[64] Danny Tang, Yuria Utsumi, and Natalie Lao. 'PIC: A personal image classification webtool for high school students'. In: *Proceedings of the International Workshop on Education in Artifial Intelligence K–12 EDUAI-20*. 2019 (cited on pages 22, 40).

[65] Michelle Carney et al. '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*. 2020, pp. 1–8 (cited on pages 22, 66).

[66] Tiffany Tseng et al. 'PlushPal: Storytelling with Interactive Plush Toys and Machine Learning'. In: *Interaction Design and Children*. 2012, pp. 236–245 (cited on pages 22, 40, 48).

[67] Abigail Zimmermann-Niefield et al. 'Youth learning machine learning through building models of athletic moves'. In: *Proceedings of the 18th ACM International Conference on Interaction Design and Children*. 2019, pp. 121–132 (cited on pages 22, 35, 40).

[68] Stefania Druga et al. 'Inclusive AI literacy for kids around the world'. In: *Proceedings of FabLearn 2019*. 2019, pp. 104–111 (cited on pages 22, 35, 40).

[69] Brian Jordan et al. 'PoseBlocks: A Toolkit for Creating (and Dancing) with AI'. In: *The 11th Symposium on Educational Advances in Artificial Intelligence*. 2021 (cited on pages 22, 27, 35, 36, 40, 66).

[70] Juan David Rodríguez García et al. 'Learningml: A tool to foster computational thinking skills through practical artificial intelligence projects'. In: *Revista de Educación a Distancia (RED)* 20.63 (2020) (cited on page 22).

[71] Tejal Reddy, Randi Williams, and Cynthia Breazeal. 'Text Classification for AI Education.' In: *SIGCSE*. New York, NY, USA: Association for Computing Machinery, 2021, p. 1381 (cited on pages 22, 27, 66).

[72] Jeffrey P Bigham et al. 'Inspiring blind high school students to pursue computer science with instant messaging chatbots'. In: *Proceedings of the 39th SIGCSE technical symposium on Computer science education*. 2008, pp. 449–453 (cited on pages 22, 40).

[73] Joy Buolamwini and Timnit Gebru. 'Gender shades: Intersectional accuracy disparities in commercial gender classification'. In: *Conference on fairness, accountability and transparency*. PMLR. 2018, pp. 77–91 (cited on page 23).

[74] National Academies of Sciences, Engineering, and Medicine and others. 'Cultivating interest and competencies in computing: Authentic experiences and design factors'. In: (2021) (cited on page 23).

[75] Gloria Ladson-Billings. 'Culturally relevant pedagogy 2.0: aka the remix'. In: *Harvard educational review* 84.1 (2014), pp. 74–84 (cited on page 23).

[76] Clifford H Lee and Elisabeth Soep. 'None but ourselves can free our minds: Critical computational literacy as a pedagogy of resistance'. In: *Equity & Excellence in Education* 49.4 (2016), pp. 480–492 (cited on pages 23, 24).

[77] Brenda A Allen and Lisa Butler. 'The effects of music and movement opportunity on the analogical reasoning performance of African American and White school children: A preliminary study'. In: *Journal of Black Psychology* 22.3 (1996), pp. 316–328 (cited on page 23).

[78] Geneva Gay. *Culturally responsive teaching: Theory, research, and practice*. teachers college press, 2018 (cited on page 23).

[79] Federico R Waitoller and Kathleen A King Thorius. 'Cross-pollinating culturally sustaining pedagogy and universal design for learning: Toward an inclusive pedagogy that accounts for dis/ability'. In: *Harvard Educational Review* 86.3 (2016), pp. 366–389 (cited on page 23).

[80] A Wade Boykin. 'Harvesting talent and culture: African American children and educational reform'. In: *Schools and students at risk* (1994), pp. 116–138 (cited on page 23).

[81] Cherry A McGee Banks and James A Banks. 'Equity pedagogy: An essential component of multicultural education'. In: *Theory into practice* (1995), pp. 152–158 (cited on page 23).

[82] Raechel Walker, Eman Sherif, and Cynthia Breazeal. 'Liberatory Computing Education for African American Students'. In: *2022 Conference on Research in Equitable and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*. IEEE. 2022, pp. 1–5 (cited on pages 23, 24).

[83] Clifford Lee and Elisabeth Soep. *Code for What?: Computer Science for Storytelling and Social Justice*. MIT Press, 2023 (cited on page 24).

[84] Jane Margolis et al. 'Beyond access: Broadening participation in high school computer science'. In: *ACM Inroads* 3.4 (2012), pp. 72–78 (cited on pages 24, 127).

[85] Jean J Ryoo. 'Pedagogy that supports computer science for all'. In: *ACM Transactions on Computing Education (TOCE)* 19.4 (2019), pp. 1–23 (cited on page 24).

[86] Jayne Everson, F Megumi Kivuva, and Amy J Ko. '"A Key to Reducing Inequities in Like, AI, is by Reducing Inequities Everywhere First" Emerging Critical Consciousness in a Co-Constructed Secondary CS Classroom'. In: *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education-Volume 1*. 2022, pp. 209–215 (cited on page 24).

[87] Jaemarie Solyst et al. '"I Would Like to Design": Black Girls Analyzing and Ideating Fair and Accountable AI'. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–14 (cited on pages 24, 38).

[88] Yinmiao Li et al. '"I Want to Be Unique From Other Robots": Positioning Girls as Co-creators of Social Robots in Culturally-Responsive Computing Education'. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–14 (cited on page 24).

[89] Jessica Morales-Chicas et al. 'Computing with relevance and purpose: A review of culturally relevant education in computing.' In: *International Journal of Multicultural Education* 21.1 (2019), pp. 125–155 (cited on page 25).

[90] Amy J Ko et al. 'It is time for more critical CS education'. In: *Communications of the ACM* 63.11 (2020), pp. 31–33 (cited on page 25).

[91] Sepehr Vakil. 'Ethics, identity, and political vision: Toward a justice-centered approach to equity in computer science education'. In: *Harvard educational review* 88.1 (2018), pp. 26–52 (cited on page 25).

[92] Casey Fiesler, Natalie Garrett, and Nathan Beard. 'What do we teach when we teach tech ethics? A syllabi analysis'. In: *Proceedings of the 51st ACM technical symposium on computer science education*. 2020, pp. 289–295 (cited on pages 25, 30, 32, 39).

[93] Michael Skirpan et al. '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*. ACM. 2018, pp. 940–945 (cited on pages 25, 30).

[94] Rua M Williams et al. 'Oh No, Not Another Trolley! On the Need for a Co-Liberative Consciousness in CS Pedagogy'. In: *IEEE Transactions on Technology and Society* (2021) (cited on page 25).

[95] Natalie Garrett, Nathan Beard, and Casey Fiesler. 'More than "If Time Allows" the role of ethics in AI education'. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 2020, pp. 272–278 (cited on page 25).

[96] Edmond Awad et al. 'The moral machine experiment'. In: *Nature* 563.7729 (2018), pp. 59–64 (cited on page 25).

[97] Abby Everett Jaques. 'Why the moral machine is a monster'. In: *University of Miami School of Law* 10 (2019), pp. 1–10 (cited on page 25).

[98] JP Lassnig. 'Development of a curriculum to teach basics of artificial intelligence'. PhD thesis. Master's thesis, Graz University of Technology, Graz, Austria, 2018 (cited on page 25).

[99] Simone Opel et al. 'Entwicklung und Reflexion einer Unterrichtssequenz zum Maschinellen Lernen als Aspekt von Data Science in der Sekundarstufe II'. In: *Informatik für alle* (2019) (cited on page 25).

[100] Daniella DiPaola, Blakeley H Payne, and Cynthia Breazeal. 'Decoding design agendas: an ethical design activity for middle school students'. In: *Proceedings of the Interaction Design and Children Conference*. 2020, pp. 1–10 (cited on pages 25, 36, 40, 43, 48, 94).

[101] Xieling Chen et al. 'Two decades of artificial intelligence in education'. In: *Educational Technology & Society* 25.1 (2022), pp. 28–47 (cited on page 26).

[102] Justin Reich. *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press, 2020 (cited on pages 26, 27, 128).

[103] Safinah Ali et al. 'Children as creators, thinkers and citizens in an AI-driven future'. In: *Computers and Education: Artificial Intelligence* 2 (2021), p. 100040 (cited on pages 26, 27, 40, 103).

[104] Tony Belpaeme et al. 'Social robots for education: A review'. In: *Science robotics* 3.21 (2018) (cited on page 26).

[105] Goren Gordon, Cynthia Breazeal, and Susan Engel. 'Can children catch curiosity from a social robot?' In: *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM. 2015, pp. 91–98 (cited on pages 26, 103).

[106] Hae Won Park et al. 'Growing growth mindset with a social robot peer'. In: *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*. IEEE. 2017, pp. 137–145 (cited on pages 26, 103).

[107] Malte F Jung et al. 'Participatory materials: having a reflective conversation with an artifact in the making'. In: *Proceedings of the 2014 conference on Designing interactive systems*. 2014, pp. 25–34 (cited on pages 26, 103).

[108] Peter H Kahn et al. 'Human creativity can be facilitated through interacting with a social robot'. In: *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. Piscataway, NJ, USA: IEEE Press, 2016, pp. 173–180 (cited on pages 27, 103).

[109] Randi Williams and Cynthia Breazeal. 'Assessments for K-12 AI literacy: A comprehensive review'. In: *Proceedings of the 2nd Workshop on AI education in K-12*. 2023, pp. 1–14 (cited on pages 27, 45).

[110] Randi Williams, Stephen P Kaputsos, and Cynthia Breazeal. '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. 17. 2021, pp. 15678–15686 (cited on pages 27, 36, 55, 59, 63, 74).

[111] Randi Williams et al. 'Doodlebot: An Educational Robot for Creativity and AI Literacy'. In: *2024 19th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. 2024 (cited on pages 27, 55).

[112] Randi Williams et al. 'AI+ ethics curricula for middle school youth: Lessons learned from three project-based curricula'. In: *International Journal of Artificial Intelligence in Education* (2022), pp. 1–59 (cited on pages 27, 48, 55, 56, 58, 82, 84, 100).

[113] Randi Williams, Sharifa Alghowinem, and Cynthia Breazeal. 'Dr. R.O. Bott Will See You Now: Exploring AI for Wellbeing with Middle School Students'. In: *2024 AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*. AAAI. 2023 (cited on pages 27, 48, 55, 84).

[114] Safinah Ali et al. 'Constructionism, ethics, and creativity: Developing primary and middle school artificial intelligence education'. In: *International workshop on education in artificial intelligence k-12 (eduai'19)*. Vol. 2. 2019, pp. 1–4 (cited on pages 29, 55).

[115] Joyce Mahon, Brett A Becker, and Brian Mac Namee. 'AI and ML in School Level Computing Education: Who, What and Where?' In: *Irish Conference on Artificial Intelligence and Cognitive Science*. Springer. 2022, pp. 201–213 (cited on page 31).

[116] Nicole M Joseph et al. 'The role of socialization in shaping Black girls' mathematics identity: An analysis of the high school longitudinal study 2009'. In: *Teachers College Record* 122.11 (2020), pp. 1–34 (cited on page 31).

[117] Remy Dou et al. 'Early informal STEM experiences and STEM identity: The importance of talking science'. In: *Science Education* 103.3 (2019), pp. 623–637 (cited on page 31).

[118] Viviane Seyranian et al. 'The longitudinal effects of STEM identity and gender on flourishing and achievement in college physics'. In: *International journal of STEM education* 5.1 (2018), pp. 1–14 (cited on page 31).

[119] Thilo Hagendorff. 'The ethics of AI ethics: An evaluation of guidelines'. In: *Minds and machines* 30.1 (2020), pp. 99–120 (cited on page 32).

[120] Claes Wohlin. 'Guidelines for snowballing in systematic literature studies and a replication in software engineering'. In: *Proceedings of the 18th international conference on evaluation and assessment in software engineering*. 2014, pp. 1–10 (cited on page 32).

[121] Gail Rosen, Jason Silverman, and Steve Essinger. 'Inquiry-based learning through image processing [dsp education]'. In: *IEEE Signal Processing Magazine* 29.1 (2011), pp. 164–169 (cited on page 40).

[122] Luciana Benotti, María Cecilia Martínez, and Fernando Schapachnik. 'Engaging high school students using chatbots'. In: *Proceedings of the 2014 conference on Innovation & technology in computer science education*. 2014, pp. 63–68 (cited on page 40).

[123] Martin Kandlhofer et al. 'Artificial intelligence and computer science in education: From kindergarten to university'. In: *2016 IEEE Frontiers in Education Conference (FIE)*. IEEE. 2016, pp. 1–9 (cited on page 40).

[124] Harald Burgsteiner, Martin Kandlhofer, and Gerald Steinbauer. 'Irobot: Teaching the basics of artificial intelligence in high schools'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 30. 1. 2016, pp. 4126–4127 (cited on pages 35, 40).

[125] Marie E Vachovsky et al. 'Toward more gender diversity in CS through an artificial intelligence summer program for high school girls'. In: *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*. 2016, pp. 303–308 (cited on pages 36, 40).

[126] Shashank Srikant and Varun Aggarwal. 'Introducing data science to school kids'. In: *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education*. 2017, pp. 561–566 (cited on page 40).

[127] Tom Hitron et al. 'Introducing children to machine learning concepts through hands-on experience'. In: *Proceedings of the 17th ACM Conference on Interaction Design and Children*. 2018, pp. 563–568 (cited on page 40).

[128] Bawornsak Sakulkueakulsuk et al. 'Kids making AI: Integrating machine learning, gamification, and social context in STEM education'. In: *2018 IEEE international conference on teaching, assessment, and learning for engineering (TALE)*. IEEE. 2018, pp. 1005–1010 (cited on pages 40, 48).

[129] K Megasari Kahn et al. 'AI programming by children using snap! block programming in a developing country'. In: (2018) (cited on pages 35, 36, 40).

[130] Jennifer Ureta and Joanna Pauline Rivera. 'Using Chatbots to Teach STEM Related Research Concepts to High School Students'. In: *Proceedings of the 26th International Conference on Computers in Education*. 2018 (cited on page 40).

[131] Tom Hitron et al. '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*. 2019, pp. 1–11 (cited on pages 35, 40).

[132] Randi Williams, Hae Won Park, and Cynthia Breazeal. '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*. 2019, pp. 1–11 (cited on pages 35, 40).

[133] Randi Williams et al. 'Popbots: Designing an artificial intelligence curriculum for early childhood education'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01. 2019, pp. 9729–9736 (cited on pages 35, 40).

[134] Julian Estevez, Gorka Garate, and Manuel Graña. 'Gentle introduction to artificial intelligence for high-school students using scratch'. In: *IEEE access* 7 (2019), pp. 179027–179036 (cited on pages 35, 40).

[135] Radu Mariescu-Istodor and Ilkka Jormanainen. 'Machine learning for high school students'. In: *Proceedings of the 19th Koli Calling International Conference on Computing Education Research*. 2019, pp. 1–9 (cited on pages 35, 40).

[136] Bamshad Mobasher et al. 'Data Science Summer Academy for Chicago Public School Students'. In: *ACM SIGKDD Explorations Newsletter* 21.1 (2019), pp. 49–52 (cited on pages 35, 36, 40).

[137] Yuanlin Zhang et al. 'A preliminary report of integrating science and computing teaching using logic programming'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01. 2019, pp. 9737–9744 (cited on pages 36, 40, 48).

[138] Nora Alturayeif, Nouf Alturaief, and Zainab Alhathloul. 'DeepScratch: scratch programming language extension for deep learning education'. In: *International Journal of Advanced Computer Science and Applications* 11.7 (2020), pp. 642–650 (cited on pages 35, 40).

[139] Karl-Emil Kjær Bilstrup, Magnus H Kaspersen, and Marianne Graves Petersen. 'Staging reflections on ethical dilemmas in machine learning: A card-based design workshop for high school students'. In: *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 2020, pp. 1211–1222 (cited on page 40).

[140] Phoebe Lin et al. 'Zhorai: Designing a conversational agent for children to explore machine learning concepts'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 09. 2020, pp. 13381–13388 (cited on page 40).

[141] Narges Norouzi, Snigdha Chaturvedi, and Matthew Rutledge. 'Lessons learned from teaching machine learning and natural language processing to high school students'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 09. 2020, pp. 13397–13403 (cited on pages 40, 85).

[142] Alpay Sabuncuoglu. 'Designing one year curriculum to teach artificial intelligence for middle school'. In: *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*. 2020, pp. 96–102 (cited on pages 36, 40).

[143] Marie-Monique Schaper, Laura Malinverni, and Cristina Valero. 'Robot Presidents: Who should rule the world? Teaching Critical Thinking in AI through Reflections upon Food Traditions'. In: *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 2020, pp. 1–4 (cited on page 40).

[144] Gilad Shamir and Ilya Levin. 'Transformations of computational thinking practices in elementary school on the base of artificial intelligence technologies'. In: *Proceedings of EDULEARN20 Conference*. Vol. 6. 2020 (cited on page 40).

[145] Zoe Skinner, Stacey Brown, and Greg Walsh. 'Children of Color's Perceptions of Fairness in AI: An Exploration of Equitable and Inclusive Co-Design'. In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–8 (cited on page 40).

[146] Jessica Van Brummelen et al. 'CONVO: What does conversational programming need?' In: *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE. 2020, pp. 1–5 (cited on pages 36, 40).

[147] Henriikka Vartiainen, Matti Tedre, and Teemu Valtonen. 'Learning machine learning with very young children: Who is teaching whom?' In: *International journal of child-computer interaction* 25 (2020), p. 100182 (cited on pages 35, 40).

[148] Henriikka Vartiainen et al. 'Machine learning for middle-schoolers: Children as designers of machine-learning apps'. In: *2020 IEEE Frontiers in Education Conference (FIE)*. IEEE. 2020, pp. 1–9 (cited on pages 36, 40).

[149] Xiaoyu Wan et al. 'SmileyCluster: supporting accessible machine learning in K-12 scientific discovery'. In: *Proceedings of the Interaction Design and Children Conference*. 2020, pp. 23–35 (cited on page 40).

[150] Eunsun Choi and Namje Park. 'Demonstration of Gamification in Education for Understanding Artificial Intelligence Principles at Elementary School Level.' In: *Ilkogretim Online* 20.3 (2021) (cited on page 40).

[151] Stefania Druga and Amy J Ko. 'How do children's perceptions of machine intelligence change when training and coding smart programs?' In: *Interaction Design and Children*. 2021, pp. 49–61 (cited on page 40).

[152] Stacey Forsyth et al. 'Imagine a More Ethical AI: Using Stories to Develop Teens' Awareness and Understanding of Artificial Intelligence and its Societal Impacts'. In: (2021) (cited on pages 36, 40).

[153] Julie Henry, Alyson Hernalesteen, and Anne-Sophie Collard. 'Teaching Artificial Intelligence to K-12 Through a Role-Playing Game Questioning the Intelligence Concept'. In: *KI-Künstliche Intelligenz* 35.2 (2021), pp. 171–179 (cited on pages 36, 40).

[154] Magnus Høholt Kaspersen, Karl-Emil Kjær Bilstrup, and Marianne Graves Petersen. 'The Machine Learning Machine: A Tangible User Interface for Teaching Machine Learning'. In: *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 2021, pp. 1–12 (cited on page 40).

[155] Pei-Yi Lin et al. 'Modeling the structural relationship among primary students' motivation to learn artificial intelligence'. In: *Computers and Education: Artificial Intelligence* 2 (2021), p. 100006 (cited on page 40).

[156] Magnus Hoeholt Kaspersen et al. 'VotestratesML: A High School Learning Tool for Exploring Machine Learning and its Societal Implications'. In: *FabLearn Europe/MakeEd 2021-An International Conference on Computing, Design and Making in Education*. 2021, pp. 1–10 (cited on page 40).

[157] Seonghun Kim et al. 'Why and what to teach: AI curriculum for elementary school'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 17. 2021, pp. 15569–15576 (cited on page 40).

[158] Duri Long et al. 'The Role of Collaboration, Creativity, and Embodiment in AI Learning Experiences'. In: *Creativity and Cognition*. 2021, pp. 1–10 (cited on page 40).

[159] Zhuoyue Lyu, Safinah Ali, and Cynthia Breazeal. 'Introducing Variational Autoencoders to High School Students'. In: *arXiv preprint arXiv:2111.07036* (2021) (cited on page 40).

[160] Gaspar Isaac Melsión et al. 'Using Explainability to Help Children UnderstandGender Bias in AI'. In: *Interaction Design and Children*. 2021, pp. 87–99 (cited on pages 36, 40).

[161] Viktoriya Olari, Kostadin Cvejoski, and Øyvind Eide. 'Introduction to Machine Learning with Robots and Playful Learning'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 17. 2021, pp. 15630–15639 (cited on page 40).

[162] Kyungjin Park et al. 'Designing a Visual Interface for Elementary Students to Formulate AI Planning Tasks'. In: *2021 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE Computer Society. 2021, pp. 1–9 (cited on pages 35, 40).

[163] Juan David Rodríguez-García et al. 'Evaluation of an Online Intervention to Teach Artificial Intelligence with LearningML to 10-16-Year-Old Students'. In: *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 2021, pp. 177–183 (cited on pages 40, 45, 85).

[164] Gilad Shamir and Ilya Levin. 'Neural Network Construction Practices in Elementary School'. In: *KI-Künstliche Intelligenz* (2021), pp. 1–9 (cited on pages 40, 48).

[165] Mariana Aki Tamashiro et al. 'Introducing Teenagers to Machine Learning through Design Fiction: An Exploratory Case Study'. In: *Interaction Design and Children*. 2021, pp. 471–475 (cited on pages 36, 40).

[166] Jessica Van Brummelen, Viktoriya Tabunshchyk, and Tommy Heng. '" Alexa, Can I Program You?": Student Perceptions of Conversational Artificial Intelligence Before and After Programming Alexa'. In: *arXiv preprint arXiv:2102.01367* (2021) (cited on pages 35, 36, 40).

[167] Henriikka Vartiainen et al. ‘Machine learning for middle schoolers: Learning through data-driven design’. In: *International Journal of Child-Computer Interaction* 29 (2021), p. 100281 (cited on pages 35, 40).

[168] Iro Voulgari et al. ‘Learn to Machine Learn: Designing a Game Based Approach for Teaching Machine Learning to Primary and Secondary Education Students’. In: *Interaction Design and Children*. 2021, pp. 593–598 (cited on pages 36, 40).

[169] Randi Williams. ‘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*. 2021, pp. 1382–1382 (cited on page 40).

[170] Spencer Yoder et al. ‘Gaining Insight into Effective Teaching of AI Problem-Solving Through CSEDM: A Case Study’. In: (2021) (cited on page 40).

[171] Ziyi Zhang et al. ‘An Interactive Robot Platform for Introducing Reinforcement Learning to K-12 Students’. In: *International Conference on Robotics in Education (RiE)*. Springer. 2021, pp. 288–301 (cited on page 40).

[172] Jessica Zhu and Jessica Van Brummelen. ‘Teaching students about conversational ai using convo, a conversational programming agent’. In: *2021 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE. 2021, pp. 1–5 (cited on pages 36, 40).

[173] Christiane Gresse Von Wangenheim, Lívia S Marques, and Jean CR Hauck. ‘Machine Learning for All–Introducing Machine Learning in K-12’. In: (2020) (cited on pages 35, 48, 58).

[174] Heloisa Candello, Mauro Carlos Pichiliani, and Claudio Santos Pinhanez. ‘Teaching Conversational Robots in a Museum Exhibition with Interactive Surfaces’. In: *Interactive Surfaces and Spaces*. 2021, pp. 41–45 (cited on page 36).

[175] Duri Long, Jonathan Moon, and Brian Magerko. ‘Unplugged assignments for K-12 AI education’. In: *AI Matters* 7.1 (2021), pp. 10–12 (cited on page 36).

[176] Cloneria Nyambali Jatileni et al. ‘Artificial intelligence in compulsory level of education: perspectives from Namibian in-service teachers’. In: *Educ. Inf. Technol.* (2023) (cited on page 37).

[177] Johanna Velander et al. ‘Artificial Intelligence in K-12 Education: eliciting and reflecting on Swedish teachers’ understanding of AI and its implications for teaching & learning’. In: Springer Science and Business Media LLC, 2023 (cited on page 37).

[178] Iro Voulgari et al. ‘Artificial Intelligence and Machine Learning Education and Literacy: Teacher Training for Primary and Secondary Education Teachers’. In: *Handbook of Research on Integrating ICTs in STEAM Education*. IGI Global, 2022, pp. 1–21 (cited on page 37).

[179] Irene Lee and Beatriz Perret. ‘Preparing High School Teachers to Integrate AI Methods into STEM Classrooms’. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 11. 2022, pp. 12783–12791 (cited on page 37).

[180] Joonhyeong Park et al. ‘Integrating artificial intelligence into science lessons: teachers’ experiences and views’. In: *Int. J. STEM Educ.* 10.1 (2023) (cited on page 37).

[181] Enrica Amplo and Deirdre Butler. ‘A Learning Programme Based on TPCK (Technological Pedagogical Content Knowledge), Constructionism, and Design to enhance Teacher Learning of the Key Ideas and Competencies of Artificial Intelligence (AI)’. In: *Proceedings of Society for Information Technology & Teacher Education International Conference 2023*. Ed. by Elizabeth Langran, Paula Christensen, and Jarrod Sanson. New Orleans, LA, United States: Association for the Advancement of Computing in Education (AACE), 2023, pp. 1914–1923 (cited on page 37).

[182] Karl-Emil Kjær Bilstrup et al. ‘Supporting critical data literacy in K-9 education: three principles for enriching pupils’ relationship to data’. In: *Interaction Design and Children*. 2022, pp. 225–236 (cited on pages 37, 38).

[183] Bich Hong Cu and Takayuki Fujimoto. ‘AI Education for Middle/High School Level: A Proposal of a System that Supports Teachers to Design Their AI Lessons’. In: *International Conference On Systems Engineering*. Springer. 2022, pp. 12–20 (cited on page 37).

[184] Daniella DiPaola et al. 'Make-a-Thon for Middle School AI Educators'. In: *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. New York, NY, USA: Association for Computing Machinery, 2023, pp. 305–311. DOI: 10.1145/3545945.3569743 (cited on page 37).

[185] Jingwan Tang et al. 'ML4STEM Professional Development Program: Enriching K-12 STEM Teaching with Machine Learning'. In: *International Journal of Artificial Intelligence in Education* 33.1 (2023), pp. 185–224 (cited on page 37).

[186] H Zhang, I Lee, and K Moore. 'Preparing teachers to teach artificial intelligence in classrooms: An exploratory study'. In: *Proceedings of the 17th International Conference of the Learning Sciences - ICLS 2023*. Ed. by P Blikstein et al. 2023, pp. 974–977 (cited on page 37).

[187] Code.org. *Lesson 4: Chatbots and Large Language Models (LLMs)*. https://studio.code.org/s/how-ai-works-2023/lessons/4. 2023 (cited on page 37).

[188] Day of AI. *ChatGPT in School*. https://www.dayofai.org/teacher-pages/chatgpt-in-school. 2023 (cited on page 37).

[189] Sara Polak, Gianluca Schiavo, and Massimo Zancanaro. 'Teachers' perspective on artificial intelligence education: An initial investigation'. In: *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 2022, pp. 1–7 (cited on page 38).

[190] Giorgia Gibellini, Valeria Fabretti, and Gianluca Schiavo. 'AI Education from the Educator's Perspective: Best Practices for an Inclusive AI Curriculum for Middle School'. In: *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2023. DOI: 10.1145/3544549.3585747 (cited on page 38).

[191] Ramon Mayor Martins et al. 'Machine Learning for All!—Introducing Machine Learning in Middle and High School'. In: *International Journal of Artificial Intelligence in Education* (2023), pp. 1–39 (cited on page 38).

[192] Ismaila Temitayo Sanusi et al. 'Learning machine learning with young children: exploring informal settings in an African context'. In: *Computer Science Education* (2023), pp. 1–32 (cited on page 38).

[193] Ismaila Temitayo Sanusi et al. 'Developing middle school students' understanding of machine learning in an African school'. In: *Computers and Education: Artificial Intelligence* 5 (2023), p. 100155 (cited on page 38).

[194] Siu-Cheung Kong, William Man-Yin Cheung, and Olson Tsang. 'Evaluating an artificial intelligence literacy programme for empowering and developing concepts, literacy and ethical awareness in senior secondary students'. In: *Education and Information Technologies* 28.4 (2023), pp. 4703–4724 (cited on page 38).

[195] Clifford H Lee et al. 'In the black mirror: Youth investigations into artificial intelligence'. In: *ACM Transactions on Computing Education* 22.3 (2022), pp. 1–25 (cited on page 38).

[196] Jean Salac et al. 'Funds of Knowledge Used by Adolescents of Color in Scaffolded Sensemaking around Algorithmic Fairness'. In: *Proceedings of the 2023 ACM Conference on International Computing Education Research - Volume 1*. New York, NY, USA: Association for Computing Machinery, 2023, pp. 191–205. DOI: 10.1145/3568813.3600110 (cited on page 38).

[197] Jaemarie Solyst et al. 'Investigating Girls' Perspectives and Knowledge Gaps on Ethics and Fairness in Artificial Intelligence in a Lightweight Workshop'. In: *arXiv preprint arXiv:2302.13947* (2023) (cited on page 38).

[198] Stefania Druga, Fee Lia Christoph, and Amy J Ko. 'Family as a Third Space for AI Literacies: How do children and parents learn about AI together?' In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–17 (cited on page 38).

[199] Gloria Ashiya Katuka et al. 'A Summer Camp Experience to Engage Middle School Learners in AI through Conversational App Development'. In: *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. 2023, pp. 813–819 (cited on page 38).

[200] Marie-Monique Schaper and Aurelio Ruiz Garcia. 'Three recommendations to engage At-Risk Students in Critical Reflection on Intelligent Technologies through Remote Learning'. In: *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–8 (cited on page 38).

[201] Ibrahim Oluwajoba Adisa et al. 'SPOT: A Game-Based Application for Fostering Critical Machine Learning Literacy Among Children'. In: *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference*. 2023, pp. 507–511 (cited on page 38).

[202] Golnaz Arastoopour Irgens et al. 'Characterizing children's conceptual knowledge and computational practices in a critical machine learning educational program'. In: *International Journal of Child-Computer Interaction* 34 (2022), p. 100541 (cited on page 38).

[203] Minji Kong and Lori Pollock. 'Experiences Piloting a Diversity and Inclusion in Computing Innovations Course'. In: *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. 2023, pp. 242–248 (cited on page 38).

[204] Xiao-Fan Lin et al. 'Technological support to foster students' artificial intelligence ethics: An augmented reality-based contextualized dilemma discussion approach'. In: *Computers & Education* 201 (2023), p. 104813 (cited on page 38).

[205] Merijke Coenraad. '"That's what techquity is": youth perceptions of technological and algorithmic bias'. In: *Information and Learning Sciences* 123.7/8 (2022), pp. 500–525 (cited on page 38).

[206] Noelle Brown, Koriann South, and Eliane S Wiese. 'The Shortest Path to Ethics in AI: An Integrated Assignment Where Human Concerns Guide Technical Decisions'. In: *Proceedings of the 2022 ACM Conference on International Computing Education Research-Volume 1*. 2022, pp. 344–355 (cited on page 38).

[207] Sunok Lee et al. 'Fostering Youth's Critical Thinking Competency About AI through Exhibition'. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 2023, pp. 1–22 (cited on page 38).

[208] Colin de la Higuera. 'A report about Education, Training Teachers and Learning Artificial Intelligence: Overview of key issues.' In: (2019) (cited on pages 39, 58).

[209] Yolanda Gil and Bart Selman. 'A 20-year community roadmap for artificial intelligence research in the US'. In: *arXiv preprint arXiv:1908.02624* (2019) (cited on page 42).

[210] J Torney-Purta and S Vermeer. 'Developing Citizenship Competencies from Kindergarten through Grade 12: A Background Paper for Policymakers and Educators, Education Commission of the States, Denver, Colorado, USA'. In: (2004) (cited on pages 43, 44).

[211] Astrid Carolus et al. 'MAILS–Meta AI Literacy Scale: Development and Testing of an AI Literacy Questionnaire Based on Well-Founded Competency Models and Psychological Change-and Meta-Competencies'. In: *arXiv preprint arXiv:2302.09319* (2023) (cited on pages 48, 85, 91).

[212] David Wade Chambers. 'Stereotypic images of the scientist: The draw-a-scientist test'. In: *Science education* 67.2 (1983), pp. 255–265 (cited on page 48).

[213] ISTE. 'Hands-On AI Projects for the Classroom: A Guide for Secondary Teachers'. In: (2022) (cited on page 49).

[214] University of Helsinki and MinnaLearn. 'Elements of AI: Introduction to AI'. In: (2021) (cited on page 49).

[215] Jeffrey Saltz et al. 'Integrating ethics within machine learning courses'. In: *ACM Transactions on Computing Education (TOCE)* 19.4 (2019), pp. 1–26 (cited on page 56).

[216] Yim Register and Amy J Ko. 'Learning machine learning with personal data helps stakeholders ground advocacy arguments in model mechanics'. In: *Proceedings of the 2020 ACM Conference on International Computing Education Research*. 2020, pp. 67–78 (cited on page 56).

[217] Anu Vazhayil et al. 'Focusing on teacher education to introduce AI in schools: Perspectives and illustrative findings'. In: *2019 IEEE Tenth International Conference on Technology for Education (T4E)*. IEEE. 2019, pp. 71–77 (cited on pages 58, 74, 81).

[218] Melissa L Stone et al. 'The design implementation framework: Iterative design from the lab to the classroom'. In: *End-User considerations in educational technology design*. IGI Global, 2018, pp. 76–98 (cited on page 58).

[219] National Council of Supervisors of Mathematics. 'Improving Student Achievement in Mathematics by Using Manipulatives with Classroom Instruction'. In: *Position Papers: Improving Student Achievement* (1997) (cited on page 63).

[220] Amruth N Kumar. 'Three years of using robots in an artificial intelligence course: lessons learned'. In: *Journal on Educational Resources in Computing (JERIC)* 4.3 (2004), p. 2 (cited on page 63).

[221] Marja-Ilona Koski, Jaakko Kurhila, and Tomi A Pasanen. '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*. ACM. 2008, pp. 32–40 (cited on page 63).

[222] Paul Talaga and Jae C Oh. 'Combining AIMA and LEGO mindstorms in an artificial intelligence coursetobuild realworldrobots'. In: *Journal of Computing Sciences in Colleges* 24.3 (2009), pp. 56–64 (cited on page 63).

[223] Blakeley H. Payne. 'Bringing design justice to the classroom'. In: (2021). Ed. by Tang Institute (cited on page 67).

[224] Joanna Goode and Jean J Ryoo. 'Teacher Knowledge for Inclusive Computing Learning'. In: (2019) (cited on page 82).

[225] Tejal Reddy, Randi Williams, and Cynthia Breazeal. 'LevelUp–Automatic Assessment of Block-Based Machine Learning Projects for AI Education'. In: *2022 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE. 2022, pp. 1–8 (cited on pages 85, 103, 122).

[226] Helen Zhang et al. 'Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study'. In: *International Journal of Artificial Intelligence in Education* 33.2 (2023), pp. 290–324 (cited on pages 85, 91).

[227] Randi Williams. 'Constructionism, Ethics, and Creativity: Developing Tools for the Future of Education with AI'. In: *2022 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE. 2022, pp. 1–3 (cited on pages 102, 103).

[228] Ruben Mitnik, Miguel Nussbaum, and Alvaro Soto. 'An autonomous educational mobile robot mediator'. In: *Autonomous Robots* 25.4 (2008), pp. 367–382 (cited on page 103).

[229] Safinah Ali, Vishesh Kumar, and Cynthia Breazeal. 'AI Audit: A Card Game to Reflect on Everyday AI Systems'. In: *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2023. DOI: 10.1609/aaai.v37i13.26897 (cited on page 107).

[230] Virginia Dignum and Klara Pigmans. *Tools to operationalize the UNICEF policy guidance on AI for children*. https://www.unicef.org/globalinsight/media/1166/file/UNICEF-Global-Insight-tools-to-operationalize-AI-policy-guidance-2020.pdf. 2020 (cited on page 107).

[231] David Anderson et al. *Ethics and algorithms toolkit*. https://ethicstoolkit.ai/. 2021 (cited on page 107).

[232] Josh A. Goldstein et al. *Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations*. 2023. URL: https://arxiv.org/abs/2301.04246 (cited on page 123).