

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING: UNPACKING HIGH
SCHOOL CS TEACHERS' PERSPECTIVES AND PEDAGOGICAL APPROACHES

By

Zachary Opps

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Educational Psychology and Educational Technology—Doctor of Philosophy

2024

ABSTRACT

As the use of artificial intelligence (AI), especially machine learning (ML), has dramatically increased, K-12 schools have begun to deliver AI education; however, little is known about teachers' views on the field. This qualitative study investigated how U.S. high school computer science (CS) teachers conceptualize AI, the role of AI in their CS instruction, and the instructional curricula and pedagogies these educators use to bring AI into their CS instruction. Data was collected through semi-structured interviews with 23 educators teaching 9-12th grade CS courses in U.S. schools. Data collected from interviews was examined using thematic analysis resulting in seven themes. The findings suggest that teachers see great value in AI education and encourage schools to provide AI literacy instruction for all students while providing AI technical instruction in elective CS courses. Teachers demonstrated high levels of interest in the field but a shallow understanding of AI technology. The study's findings also showed that CS teachers know of the importance of AI ethics instruction but have a limited view of AI's impact on society. Additionally, the results point to a general lack of curricula and tools designed to teach K-12 students about AI, especially materials that emphasize critique of AI technology and its societal harms. The study's results contribute to a deeper understanding of 9-12th grade teachers' conceptions of AI and the challenges they face when implementing AI instruction. Implications for teachers, school leaders, curriculum developers, policy makers, teacher educators, and professional development (PD) providers are presented.

ACKNOWLEDGEMENTS

This study has helped me understand how collaborative any good research is. Thank you to my advisor and dissertation director, Dr. Aman Yadav, for creating the opportunity for me to conduct this research at Michigan State University, providing gentle pushes out of my comfort zone, challenging me to extend my thinking, and offering guidance in navigating the inevitable wrinkles that arise in the research process. I am deeply grateful to my dissertation committee members: Dr. Christine Greenhow, Dr. Kristen Johnson, and Dr. Anne Leftwich. I would like to extend my gratitude for their insightful feedback and suggestions that have refined this work. Also, this document would not have been possible without the many people who modeled curiosity and provided unrelenting support, patience, encouragement, and advice. Many thanks. Finally, I would like to express special thanks to the teachers who participated in the study not only for their shrewd insights but more importantly for the invaluable work they do daily with their students.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION.....	1
CHAPTER 2: LITERATURE REVIEW.....	8
CHAPTER 3: METHODS.....	51
CHAPTER 4: RESULTS.....	67
CHAPTER 5: DISCUSSION.....	128
REFERENCES.....	147
APPENDIX A: RECRUITMENT MESSAGE.....	167
APPENDIX B: PARTICIPANT SCREENING SURVEY.....	168
APPENDIX C: CONSENT FORM.....	170
APPENDIX D: INTERVIEW PROTOCOL.....	172
APPENDIX E: INTERVIEW UNPLUGGED LESSON PRESENTATION.....	175
APPENDIX F: CODEBOOK.....	179

CHAPTER 1: INTRODUCTION

Background

Artificial Intelligence (AI) and Machine Learning (ML) software systems are permeating our world. These rapidly progressing technologies are becoming more commonplace throughout business and industry and society as a whole. AI have become prevalent in a variety of frequently used products and services including smart digital voice assistants, self-driving features in vehicles, smart home devices, entertainment suggestion algorithms, mapping and travel directions software, automated investing in financial markets, online shopping recommendations, search engine algorithms, facial recognition in electronic devices, writing and grammar suggestions in text editing software, and social media feed personalization. Their widespread adoption is dramatically impacting how we work, the type of work we do, and how we interact with one another. The projected near ubiquity of AI offers the potential to further automate menial tasks currently performed by humans, such as chatbots providing customer support rather than human customer service agents or responsive manufacturing robots that work together to assemble goods rather than human assembly workers. If implemented thoughtfully, these advances promise to free up humans to engage in more meaningful work, increase our productivity, and improve our standard of living (Castro & New, 2016; Makridakis, 2017).

AI's wide-ranging functionality allows it to be deployed in a variety of industries where its capabilities can swiftly alter the way work is conceived and performed. Healthcare, manufacturing, transportation, foodservice, entertainment, agriculture, finance, education, and many other sectors of the economy have implemented AI

technology to improve workflows. These enhancements have been found to streamline industries' operations and cut time and effort spent on monotonous duties. AI has also dramatically reshaped careers and even eliminated the need for certain long-standing professions that can now be automated (Frey & Osborne, 2017). At the same time, AI systems have brought rise to new jobs that design, build, and make use of AI technology in order to create these efficiencies (Zahidi et al., 2020).

However, AI's seemingly boundless promise is not without significant dangers. Experts warn of threats to privacy, elections, information reliability, decision-making transparency, and equality (Manheim & Kaplan, 2019; Rozado, 2023). AI has also been increasingly woven into the decision-making process of key structures in the United States (U.S.). As the software becomes further embedded into the criminal justice (Završnik, 2020), education (Murphy, 2019), employment (Raub, 2018), housing (Zou & Khern-am-nuai, 2022), financial (Johnson et al., 2019), and other structures, vulnerable populations' equitable access to and unprejudiced treatment from these essential structures could be further threatened.

Should there be increased government oversight and regulation to ensure that the application of AI in specific technologies protects human rights and does not infringe on civil liberties (Bradley et al., 2020; Koepke et al., 2023)? Instead of placing blind trust in technology visionaries who develop AI and corporations who employ the technology, should consumers always be made aware when they are interacting with AI systems and be provided with a transparent explanation of how the technology makes its decisions (Bhatt et al., 2020; Doshi-Velez et al., 2017)? Does the population deserve an AI Bill of Rights of some kind? In the U.S., the debates over the role the government

should play in providing protections for its citizens generally take place through public discourse and dialogue with government officials. In order to prepare citizens to engage in this discourse, there is a need to grow a foundational understanding of AI that can democratize the technology by developing the public's understanding of AI's role, impacts, and oversight.

As the world increasingly moves toward automation, it becomes more essential for the citizenry to co-exist with AI technology to succeed in the contemporary workplace. Additionally, educating society about how AI technologies are designed and operate might work to empower citizens to understand and overcome the prejudice, discrimination, and harms that can be perpetuated by these systems. Furthermore, computing technologies are often viewed by users as evenhanded and beyond reproach, and a deeper understanding of the technology is necessary to help citizens challenge these conceptions (O'Neil, 2016). Informed citizens are better able to consider how these systems have the capacity to both protect and infringe on their rights and are also empowered to engage in dialogue about the benefits and drawbacks of employing AI in crucial areas of society. Additionally, a cognizant public can maintain a healthy skepticism, ask revealing questions, and champion for appropriate regulation and controls amidst the hype about the revolutionary systems being deployed. However, this consideration can only happen with the requisite comprehension of AI for everyone (Long & Magerko, 2020; Ng et al., 2021a; Ng et al., 2021b).

One logical place to start developing informed citizens is in K-12 schools. Providing students with exposure to AI will allow them to make sense of, contribute to, and shape an AI-filled future. Today's students will become tomorrow's decision makers.

As such, awareness of how AI can be misused could allow them to identify and remedy intentional and inadvertent misapplication of this potent technology. Given the need to create conscious consumers of AI, engaged citizens who can evaluate the technology's impact on the public, and future workers prepared to contribute to an economy infused with AI, it is imperative that students learn about these systems. Many U.S. secondary schools currently teach the closely-related subject computer science (CS) through standalone courses or by incorporating CS concepts into other core subjects such as science, math, and language arts. A logical extension to this CS instruction is to expand these curricula to include associated AI concepts. Competence in how the systems function could empower CS students who may become future software developers to identify potential ethical lapses and sources of bias at all stages in the design and development of AI.

The content taught in U.S. K-12 schools is largely dictated by state-wide content standards that typically spell out what concepts should be taught at which grade level. In an effort to align standards across the country, many states often adopt a version of a uniform set of standards such as the Common Core State Standards for Mathematics and Language Arts (National Governors Association Center for Best Practices, Council of Chief State School Officers, 2010) or the Next Generation Science Standards (NGSS Lead States, 2013). To guide the content taught in their schools' CS courses, many states have adopted the Computer Science Teachers Association's (CSTA) CS Standards, sometimes with slight modifications (Computer Science Teachers Association, 2017). AI is addressed twice in the current version of the CSTA standards

and both AI-related standards are at the 9-12th grade levels, suggesting that the high school level is a logical place to at least initially embed AI instruction.

If AI should be taught in high schools, questions abound about where, when, and how students should gain this experience and knowledge. What knowledge and skills do 9-12th grade students need in order to thrive in the AI era? Is it crucial that every single student learn about AI or should the focus be on those who express interest in the field? Should the goal of AI instruction be to help students build literacy, the ability to recognize and understand AI conceptually, or should the goal be to empower students with an understanding of the technical underpinnings and ability to write AI systems? When and how should students be prepared to consider the questions about responsible use and equitable development of AI technology? How can students prepare to engage in life-long dialogue about how we ensure that AI serves all members of society? The answers to these questions will best be shaped through empirical research. As educators and school systems rush to meet the growing demand for AI instruction to keep pace with its rapid development, it is critical that sufficient research is available to inform and shape their decision-making in this nascent field. However, research in the area is in its early stages and little is currently known about 9-12 AI education. One obvious place might be to begin AI education in high school CS courses, but there is little understanding of CS teachers' knowledge of AI and their views on AI in their courses.

Study Purpose

The purpose of this study was to examine AI education through the eyes of 9-12th grade U.S. CS teachers. It sought to better understand how CS teachers think about AI and the optimal means for providing AI instruction in their CS courses. Further, it explored the challenges and barriers teachers face as well as the teaching resources and strategies they use as they look to implement AI instruction in their classrooms.

Specifically, the study was guided by the following research questions:

RQ1: How do 9-12th grade CS educators conceptualize AI and ML?

RQ2: How do 9-12th grade CS educators conceptualize the role of AI in their CS instruction?

RQ3: What, if any, instructional curricula and pedagogies do 9-12th grade CS educators use to bring AI into their CS instruction?

To explore these research questions, the study used in-depth interviews to create space for participants to speak openly about their understanding of AI and ML, AI's part in their courses, and how they currently teach or would teach AI. Interviews were held with 23 CS teachers using a semi-structured interview protocol, and participants' responses were transcribed, coded, and explored through thematic analysis. Greater detail on the study's methods is provided in Chapter 3.

Overview of Dissertation

Chapter One of this dissertation introduces the study's topic, explains its purpose, provides an overview of the study, and lays out a roadmap of the dissertation document. Definitions for AI and ML, the growth and influence of the technologies, as well as the extant research literature regarding AI education and AI literacy are

discussed in Chapter Two. Chapter Three presents a positionality statement from the researcher and addresses the research methods and procedures that were followed while conducting the study. The results of the research are detailed in Chapter Four, including the codes and themes that arose from the data analysis of the interviews. These themes are explored in detail with explanations and example quotes from participant interviews. Finally, Chapter Five focuses on the findings for each research question, implications of the findings, the study's limitations, and suggested areas for further research.

CHAPTER 2: LITERATURE REVIEW

Artificial Intelligence and Machine Learning

John McCarthy and his colleagues (1955) coined the term *artificial intelligence* in a seminal research proposal. They hypothesized that a defining component of an AI system is that the machine can improve itself. That central concept lives on even as scholars offer a wide variety of definitions of the term. McCarthy subsequently defined AI as the “science and engineering of making intelligent machines, especially intelligent computer programs” and to explain that intelligence referred to the “computational part of the ability to achieve goals in the world” (McCarthy, 2007, p. 2). Alternatively, Turing (1950) posited that in order for a machine to be deemed intelligent, humans should not be able to distinguish between interactions with the machine and another human. In other words, the key is not whether a machine truly is intelligent, but whether it can persuade humans that it is intelligent.

Contemporary researchers have extended these lines of thinking with particulars framed around present technology. For example, Wang (1995, 2019) emphasized that any working definition of AI must center on the system’s versatility and flexibility which allow it to function in ill-defined circumstances that provide a limited frame of reference. Another definition by Elliott (2019) described AI as computer-based systems able to sense the context of situations and respond appropriately based on the information it is presented. In addition, Baker et al. (2019) posited that a defining characteristic of AI is that a system can engage in the type of work that has traditionally required a human brain, such as problem-solving. Subsequently, Russell and Norvig (2021) expanded on these notions to add that true AI requires a machine to have the ability to communicate

with humans (natural language processing), capture and save information (knowledge representation), apply this saved material to new problems and scenarios (automated reasoning), and adjust to new and novel situations and conditions (ML). On the other hand, the Center for Integrative Research in the Computing and Learning Sciences's AI in Education Working Group defines AI as "any computational method that is made to act independently towards a goal based on inferences from theory or patterns in data" (Friedman et al., 2021).

In 1959, Samuel introduced the machine learning approach to AI by predicting that the ML technique would allow computer programs to gain knowledge from their experiences, and reduce the need for humans to supply detailed instructions for computers (Samuel, 1959). The ML method trains intelligent agents, or algorithms, through the use of vast sample data to be able to perform a specific task, such as making predictions, without a human explicitly programming all of the instructions to do so (Mitchell, 1997; Murphy, 2012). O'Neil (2016) explained the process as follows: "The computer dives into the data, following only basic instructions. The algorithm finds patterns on its own, and then, through time, connects them with outcomes. In a sense, it learns" (p. 75). The malleability of ML technology allows for its application in a wide range of fields and settings and has led to meteoric growth in its adoption. This flexibility has led ML to become the primary AI approach, making it increasingly important for learners to understand (Sarker, 2021).

AI Growth

The use of AI has increased dramatically in recent years as evidenced by world-wide publications, including peer-reviewed journal articles, conference papers,

repositories, books, and book chapters, which have more than doubled from 162,444 in 2010 to 334,497 in 2021 (Zhang, Maslej, et al., 2022). The size of the workforce entering the AI field has seen similarly substantial gains as shown by the number of students majoring in CS. Individuals designing and implementing AI and ML systems often study CS in college before beginning work in AI professions, and the number of CS undergraduate degree earners at North American doctoral universities grew 3.5 times from fewer than 9000 in 2010 to 31,840 in 2020 (Zweben & Bizot, 2022). At the same time, global investment in AI companies has also jumped almost forty-fold from less than 2 billion U.S. Dollars (USD) in 2010 to over 77 billion USD in 2021 (VB Staff, 2021). Patents filed in the field across the globe grew almost thirty-fold from fewer than 5000 in 2015 to 141,240 in 2021 (Zhang, Maslej, et al., 2022). As a result of this dramatic growth, AI is impacting our everyday lives in the form of product recommendations on online retailers' websites, voice assistants' responses to questions through smart speakers, estimated commute times from smartphones, movie recommendations delivered by media streaming services on TVs, and a myriad of other functions. Beyond these commonplace consumer applications, the use of AI systems is also becoming increasingly widespread across industries including hiring algorithms that reduce time managers spend reading resumes; fraud detection software that attempt to flag suspicious purchases or activity; autonomous driving features that control a car if it drifts out of a lane or in emergency braking situations; and robotics algorithms used for manufacturing goods, preparing food, cleaning, weeding farm fields, and transporting raw materials.

AI were initially applied to relatively simple tasks that required little complex decision-making such as optical character recognition of text (Caulfield & Maloney, 1969) and solving mathematics word problems through arithmetic and straightforward algebra (Crevier, 1993). However, subsequent generations of the software have been designed to harness the greater processing power, storage space, and memory made available in contemporary computing hardware. As a result, these increasingly sophisticated systems can handle more advanced functions that require higher levels of decision-making. For example, these decision-making capabilities allowed experts to design AI systems to compete against the world's most skilled human players of complex strategic games such as chess (McCorduck & Cfe, 2004) and Go (Silver et al., 2016), interpret wording inferences when translating natural languages (Devlin et al., 2018), and determine how to respond appropriately to ever-changing surroundings when controlling autonomous vehicles (Yurtsever et al., 2020).

While AI software has developed greater capacity and wider application, the speed and efficiency of training the systems has also vastly improved. Between 2018 and 2021, the time required to train an image classification system dropped by 94.4% while the expense of this training fell by 63.6% (Zhang, Maslej, et al., 2022). Experts project that AI's ability will continue to progress and that the resultant systems, which will be considerably faster and more capable, will allow the adoption of the technology in a greater number of new fields and economic sectors (West & Allen, 2018; Littman et al., 2021). These tools will have the adaptability to be implemented in an even wider array of applications such as bolstering cybersecurity (Sarker et al., 2021) and providing

precision medical diagnoses and treatments (Davenport & Kalakota, 2019), but to do so the technology must have the trust of end users and buy-in of decision-makers.

As the level of adoption and technical abilities of AI systems continue to advance, popular understanding and opinion of the technology is also evolving. Even though the general public underestimates the pervasiveness of AI on the whole (Zhang & Dafoe, 2022) and as much as half of the U.S. may be unfamiliar with AI (Morning Consult/Politico, 2017), nearly 50% of the U.S. general public feel that the development of AI has been a good thing (Funk et al., 2020). Meanwhile, businesses are increasingly optimistic that they can use AI to transform their operations. Ransbotham et al. surveyed over 2500 executives in 2019 and roughly 90% of respondents, regardless of their industry, felt that AI provided a business opportunity for their organization. This widespread optimism among companies is also accompanied with a healthy fear that competitors will figure out a way to harness AI before they do and secure a competitive advantage over their organization (Ransbotham et al., 2019).

Impact on Workforce

Many processes that have previously required human judgment can now be accomplished using AI. In 2016, Ng suggested that any mental task that requires one second or less of human thought could currently or in the near future be automated using AI. Essentially, every segment of human society that has associated data will be touched by automation or process improvement AI in some way. This proliferation of AI has the potential to greatly disrupt the modern workforce. In 2018, Brynjolfsson et al. predicted that the application of AI software will impact virtually every profession in the U.S., leading to a dramatic reorganization of the work performed by humans and

machines. While their projections vary, some experts predict that nearly half of current jobs in the U.S. are at risk of being automated (Frey & Osborne, 2017). Forecasts for widespread AI application has led the World Economic Forum to project that 85 million jobs worldwide will be replaced by automation by 2025 (Zahidi et al., 2020). In a survey of nearly 13,000 workers across 18 countries in 2023 by the Boston Consulting Group, 36% of participants stated that their job is likely to be replaced by AI (Beauchene et al., 2023). As experts anticipate that 70% of businesses will employ AI by 2030, it is important that citizens learn to harness AI's capabilities in order to be competitive in the contemporary economy which is automating at an accelerating pace (Bughin et al., 2018).

While some positions and even occupations will inevitably be eliminated, other careers will likely flourish, especially those geared toward identifying automatable procedures, developing and employing the necessary AI systems to perform those tasks, and maintaining and modifying the systems to meet the needs of organizations. The World Economic Forum estimates that 97 million new jobs will be created by 2025, largely in data-, AI-, and robotics-focused fields (Zahidi et al., 2020). Historically, as a percentage of all job postings in the United States, AI and ML positions grew six-fold from below 0.15% in 2010 to 0.90% in 2021 (Zhang, Maslej, et al., 2022). This expansion is expected to continue as the U.S. Bureau of Labor Statistics projects that employment in the computer and information research science field will grow 21% between 2021 and 2031, over five times the average growth rate for all career fields in the country (Bureau of Labor Statistics, 2023). Similarly, the World Economic Forum anticipates AI and Machine Learning Specialists to be the fastest growing career

worldwide from 2023 to 2028, expanding by over 30% (Di Battista et al., 2023). This shift could present a unique and grand opportunity to upskill workers and bring economic prosperity to traditionally less skilled laborers throughout society.

However, without the necessary plans to prepare the workforce for the opportunities AI can offer, the introduction of the technology has been found to exacerbate existing economic inequalities. This has led many to question whether the transition to AI will further concentrate economic wealth to a few if a forward-looking plan is not in place to fairly distribute the potential benefits (Acemoglu & Restrepo, 2020; Brynjolfsson, 2022). However, as this great shift in the workforce accelerates, employers do not appear to be sufficiently training and retooling their personnel in order to provide opportunities for those employees uprooted by AI to pivot to newly created roles to support and leverage the technology. Of the roughly 13,000 respondents that the Boston Consulting Group surveyed world-wide, 86% stated that they will need upskilling in order to be competitive in the modern AI era yet only 14% of frontline workers, those most likely to be displaced by AI, have received any form of training to help adapt to the changes AI has brought to their profession (Beauchene et al., 2023). Lacking the proper vision, planning, and training, the economic benefits this rapid influx of new jobs could provide to underserved populations will not materialize. Economic disparity is not the only equity concern brought about by AI.

Equity and AI

New technologies lead to new problems; they can also exacerbate old ones. Wider application of AI that seeks to replicate human decision-making also threatens to replicate human biases and prejudices. Organizations world-wide increasingly use AI in

an attempt to add objectivity to their decision-making with an underlying belief that the technology is neutral (Yadav & Heath, 2022). However, AI systems are developed by people and are only as objective as the datasets they are trained with and human-written instructions they follow. The very tools that are being employed in an attempt to improve objectivity in decision-making often threaten to undermine the fairness of these important decisions (Rainie & Anderson, 2017).

Examples are becoming increasingly commonplace. Computer-based models used to evaluate applicants have been found to reflect pre-existing racial and gender bias in employment opportunities (Dastin, 2018) and medical school admissions (Lowry & Macpherson, 1988). Automated facial recognition systems have demonstrated racial and gender bias embedded into their core logic (Buolamwini & Gebru, 2018) and AI systems used to support health care providers have been shown to repeat disproven race-based misinformation (Omiye et al., 2023). Researchers have also detailed the discriminatory impact that algorithms can have on criminal justice sentencing due to racial bias built into the software (Chouldechova, 2017). For instance, recidivism prediction software such as Correctional Offender Management Profiling for Alternative Sanctions, or COMPAS, has been found to recommend significantly higher prison sentences for Black and Brown defendants compared to White defendants (Angwin et al., 2016). O'Neil (2016) illustrated how a society's discriminatory beliefs become inherent in predictive policing algorithms and the models then serve to reinforce those biases.

Suresh & Guttag (2021) identified a myriad of points throughout the typical ML life cycle that allow for the introduction of bias which contaminate the systems and lead

to negative consequences for users and innocent bystanders. Since the authors of any software are susceptible to the biases innate in their society, AI software developers' values, assumptions, and decision-making while writing the software are prone to contaminating the computer code undergirding the systems they develop (Citron & Pasquale, 2014). In addition to the influence of developers' implicit biases during the development process, AI systems' fairness can also be dramatically impacted by the datasets used to train them.

ML models are typically trained and refined using substantial collections of data points known as datasets. These datasets are frequently drawn from massive quantities of publicly available data points generated by human behavior such as dialogue from internet message boards, customer reviews, user-shared images, or buying habits. These datasets can replicate societal biases through the actions of those who create the data points that end up training the ML system (Barocas and Selbst, 2016). For example, computer models trained with datasets of news articles have shown female/male gender stereotype bias (Bolukbasi et al., 2016) while datasets of images and labels scraped from the internet can lead to systems that are trained to refer to people in racist and misogynistic ways (Quach, 2020). As inherently biased systems improve in scale and speed, AI-/ML-based tools will be positioned to magnify flawed decision-making because they will touch the lives of so many people. Thus, the resultant AI systems serve to not only reinforce but to amplify existing prejudices and inequalities of the status quo. Though many have been found to be biased, automated systems increasingly serve as gatekeepers to opportunities in our society.

Housing, employment, policing, criminal justice, political campaigns, and many other sectors of society are incorporating AI into their decision-making at an alarming rate. These automated systems can influence whether potential homeowners are approved for mortgage loans, job applicants are considered for open positions, citizens are under- or over-scrutinized by law enforcement, defendants receive impartial treatment in the criminal court system, and voters are targeted by coercive political campaigning. AI technology now plays a key role in these decisions that influence citizens' quality of life yet little government regulation exists to oversee their use. Despite the outsized influence that AI software has in our society, a mere three bills addressing some aspect of AI were passed into federal law in the United States in 2021 while only 18 federal laws were enacted world-wide in the same year (Zhang, Maslej, et al., 2022). Without proper government oversight, society runs the risk of biased AI systems jeopardizing the rights of its citizenry. Citizens who are informed about AI are in a better position to understand the impacts AI has on society as well as help shape necessary government policy and regulation.

The continued growth of AI promises to create novel situations and concerns that have not previously existed. Humans will only be able to grasp these new complexities through a conceptual, and perhaps technical understanding of how AI are designed and function. Despite the importance of being able to do so, many users struggle to recognize when software makes use of AI technology. As these systems operate in the background, the user can be unaware and essentially unknowingly accept their decision-making and recommendations. This can prevent the user from questioning how the decisions were made or challenging their fairness (Curtis et al., 2022). High

levels of use of technologies that employ AI can result in common acceptance of the systems without a critical lens towards the tools (Susser, 2019). Users' inability to recognize AI in their day-to-day lives not only limits their ability to critique AI systems but also limits citizens' ability to have informed and thorough dialogue about public policy regarding the technology (Fast & Horvitz, 2017).

In the U.S., this type of public discourse often plays out as citizens and their elected government representatives debate the merits of proposed legislation. During the 2022 session of the U.S. Congress, there were 130 AI-focused bills compared to only one in 2015 (Omaar, 2022) with some calling for the U.S. government to take legally-binding action on behalf of residents to discourage harm and encourage fairness (Blueprint for an AI bill of rights, 2022; Lander & Nelson, 2021). Approaching officials with proposals for realistic government action on AI requires a certain level of AI literacy that will allow the general public to form knowledgeable opinions about the need for protections, limits, and freedoms through government regulation, policies, or other means. AI-informed citizens can contribute in other ways as well. In a report addressing the rising threat of malicious AI use, Brundage et al. (2018) called for increased efforts to boost AI understanding in order to bring more stakeholders into conversations about risks and mitigation strategies.

AI Literacy

As the technology becomes more indispensable in society, Ng et al. (2021a) called for AI to become a "fundamental skill for everyone, not just for computer scientists. In addition to reading, writing, arithmetic and digital skills, we should add AI to every learners' twenty-first century technological literacy in work settings and everyday

life” (p. 9). The need for AI literacy is important for individuals to be able to evaluate the abilities and limits of AI systems. Ehsan et al. (2021) found that users with lower levels of AI literacy can be more easily duped into overestimating AI systems’ capabilities. For example, when presented with scenario-based decisions, medical clinicians with low levels of AI literacy picked medical treatments recommended by ML systems at a rate seven times higher than clinicians with higher AI literacy (Jacobs et al., 2021). As the use of AI technology grows, limited AI literacy could lead to humans’ overreliance on the systems’ suggestions in a growing number of critical fields, jeopardizing safety without questioning the software’s proposed solutions.

Despite the heightened need for AI literacy, mounting evidence suggests that citizens do not understand how AI systems are designed and work. In a survey of 2103 respondents based in the United Kingdom, Holder et al. (2018) found that while 75% of participants were aware of what AI is, fewer than 15% answered that they have had direct contact with AI, and only 2 in 100 believed that AI is having an impact on society. These low rates led the authors to surmise that the general public, including the segment that is familiar with AI, has limited ability to identify the technology even as they interact with it and also grossly underestimate AI’s effects on the greater community. Similarly, even though the general public is largely familiar with the existence of and the use of AI services and hardware, they rarely understand the underlying technology or comprehend the ethical concerns associated with AI systems (Burgsteiner et al., 2016). In essence, few recognize or comprehend AI even as they are becoming ubiquitous, revolutionizing many industries, changing our way of life, and intensifying societal inequities (West & Allen, 2018).

Since AI is embedded into a range of technologies children use on a regular basis, these systems serve as a looking glass that shapes how they see the world. As such, children are entitled to recognize AI, understand how the systems function, and appreciate their own potential to influence and change the technology (DiPaola et al., 2022). Many feel that children's frequent use of AI must consequently lead to a grasp of how the software works. However, mere use of AI tools alone does not remedy learners' lack of understanding of the technology (Kong et al., 2021). To combat this lack of AI literacy there have been calls to consider teaching AI at the K-12 level including in CS courses (Kandlhofer, et al., 2016; Touretzky et al., 2019). Up to this point, a limited body of research has explored how AI education is currently being implemented at the K-12 grades.

AI in K-12 Education

Increasingly, researchers have argued for AI education to be embedded in K-12 classrooms considering schools' great influence on the country's population (Roschelle et al., 2020). The general populace appears to agree. In an effort to better understand the public's views on AI in education, Orth and Bialik (2023) conducted an online survey of 1,000 U.S. adult citizens. The researchers randomly sampled opt-in participants using sample matching and weighting by gender, age, race, education, and other factors to generate a representative yet random sample. Sixty-one percent of respondents indicated that they feel it is very or somewhat necessary for K-12 students to learn AI-related skills to equip them for future careers. Teaching K-12 students about AI could have a far-reaching impact on the country's AI literacy from a very young age (Touretzky et al., 2019). A growing number of educators have recognized this fact and

K-12 AI instruction has experienced growth in the last several years (Marques, et al., 2020). However, the nascent subject does not have established best practices for K-12 instruction and content matter. The bulk of AI education research efforts have focused on university level students rather than K-12 instruction, and the limited number of K-12 studies often concentrate on one particular AI tool or concept (Burgsteiner et al., 2016). In an effort to improve on this limited, scattershot method, researchers have embarked on a variety of approaches to expand and grow AI learning opportunities for K-12 students.

Some of these efforts have been aimed at informal learning spaces, such as student competitions, summer camps, and Massive Open Online Courses (MOOCs), while others have focused on formal learning spaces, typically traditional classroom settings in schools (Wang et al., 2022). Informal AI education can offer flexibility and nimbleness unmatched by schools, allowing instruction to be responsive to new technologies, approaches, and trends. This allows informal learning spaces, which often have fewer bureaucratic requirements, to meet dynamic needs more quickly than formal learning organizations. Availability of informal learning, especially web-based virtual learning, has the potential to provide access to AI learning opportunities to learners around the world, if they seek them out. Examples of informal learning approaches to AI education include Elements of AI (EoAI), a MOOC geared toward K-12 learners, and Technovation AI Family Challenge, a competition that enlists K-12 students and their families to solve problems in their community using AI. On the other hand, while formal learning organizations are often slower to respond to needs, their systematic, methodical approach can offer increased structure and supports for school systems,

educators, and learners. The guidelines, standards, and curricula commonly found in formal education settings generally create more consistent instructional delivery and an assurance that all learners have access to the educational opportunities (Steinbauer et al., 2021). Examples of formal learning initiatives include Code.org's AI curriculum with targeted lessons geared toward 3rd to 12th grades, the International Society for Technology in Education's (ISTE) hands-on project guides aimed at elementary and secondary CS and elective courses, AI4ALL's interdisciplinary open learning lessons designed for learners without extensive CS or math backgrounds, and the AI Education Project's (aiEDU) 5-minute AI snapshot class warm-ups for 7th to 12th grade learners and series of Intro to AI project-based learning lessons aimed at 9-12th graders. Whatever the mode, both AI formal and informal learning approaches need to consider students' motivations to engage with the field.

Much is yet to be discovered about what drives students to learn about AI. Chai et al. (2021) delivered an online questionnaire to 682 students in 3rd to 6th grade in a school in Beijing, China. Their study employed the Theory of the Planned Behavior (Ajzen, 1991) to examine students' intention to learn AI. At the time of the survey, participants were enrolled in a mandatory six-hour introductory AI class. The instrument measured the following five factors: self-efficacy in learning AI, AI readiness, perceptions of the use of AI for social good, AI literacy, and behavioral intention. Using structural equation modeling, the researchers' analysis suggested that students' self-efficacy in learning AI, AI readiness, and perceived use of AI for social good were all direct predictors of students' intention to learn AI.

Regardless of the teaching approach or motivation to learn about AI, instruction requires age appropriate learning tools and instructional materials to support K-12 learners.

AI Education Tools

In a systematic review of 49 research articles, Ng, Lee, et al. (2022) found that one of the greatest barriers to K-12 AI education has historically been the shortage of beginner-appropriate software tools and accompanying teaching resources. This challenge is critical to overcome, as learners must have the ability to tinker with actual AI technology in order to comprehend the underlying concepts (Touretzky et al., 2019). Promisingly, Gresse von Wangenheim et al. (2021), in a systematic mapping of 56 platforms geared toward teaching ML, found that the development of user-friendly AI tools designed specifically for teaching and learning has gained momentum in recent years. The researchers noted that the bulk of these emerging tools were designed in a manner to allow novice learners to build AI artifacts without burdensome technical CS skill. This influx of AI education platforms has been documented by a number of studies.

Sanusi et al. (2022), in their systematic review of 39 articles from 2010 to 2021, detailed the three most common forms of instructional software tools geared toward K-12 AI learners - conversational agents, programming environments, and robotics platforms. Conversational agents, often in the form of chatbots, make use of natural language processing to interact with the learner using human language. Programming environments allow students to write programs that incorporate AI through drag-and-drop block-based or traditional programming languages such as Java, Python, R, JavaScript, or C / C++ / C#. Meanwhile, robotics tools enable students to code the

software that controls virtual or tangible robots. A primary focus for many of these AI learning platforms is to lower the barriers to entry for newcomers.

Creating AI systems from scratch generally requires a wealth of technical design skill and programming experience, which the bulk of K-12 students do not yet have. In an effort to empower beginning learners without this background to design and create AI artifacts, developers have created a variety of no- and low-code tools aimed at newcomers. The Machine Learning for Kids website, machinelearningforkids.co.uk (Lane, 2017), is one such platform. Its graphical user interface walks beginners through a process to create labeled categories and place images, text, or numbers in the appropriate bucket that exemplifies the category. The site then uses IBM Watson servers to create a classification model that can be used to classify new images based on the original categories. Next, the tool can add this newly generated model to the Scratch (Resnick et al. 2009), Python (Van Rossum, 2007), or App Inventor (Abelson et al., 2012) programming environment, allowing the learner to write programs that make use of the new image recognition rules.

Similarly, Google's Teachable Machine allows learners to build and experiment with ML classification models through a web-based graphical user interface, eliminating the need for programming expertise. Its flexibility and low barriers to entry have made the tool popular in K-12 classrooms across the globe and a number of curriculum developers have incorporated it into a variety of courses of study (Carney et al., 2020). TensorFlow Playground, also developed by Google, is a web-based, graphical tool that encourages learners to design, build, and experiment with neural networks to work out regression or classification problems without the need for coding (Smilkov et al., 2017).

Similarly, Code.org's (2023) web-based AI for Oceans tool allows learners to train a machine learning model through clicking on or tapping image training data. The graphical, experiment-friendly nature of these types of tools frees up students to focus on big picture, conceptual understanding rather than concentrating on the proper syntax necessary to make the underlying AI work properly.

Tangible tools that have been designed to illustrate AI concepts have demonstrated promising results through research with students of varying ages. For example, Druga (2018) introduced the Cognimates platform to 107 learners aged 7 to 14 years old from the U.S., Germany, Denmark, and Sweden. The software allowed users to program internet-connected toys, robots, and intelligent devices such as the Google Home and Amazon Alexa. The author found that as they programmed and played through the platform students' understanding of AI concepts deepened as well as their doubt about the devices' truthfulness and true intelligence. Williams et al. (2019) built PopBots, a hands-on platform for preschool children that connected a smartphone to LEGO pieces and motors to build a social robot which learners programmed using a block-based programming environment on a touchscreen tablet. The researchers also authored a complementary constructionist-inspired (Papert, 1980), collaborative group curriculum to teach three AI concepts - knowledge-based systems, supervised machine learning, and generative music AI through creative exploration. While delivering the curriculum as a week-long module to 80 pre-kindergarten and kindergarten students, the authors found that even young students were able to comprehend AI concepts.

Developers have also created a variety of other tools with potential for K-12 AI education. For example, the Cozmo (Brehm et al., 2019) and Create (Fernandes, 2016)

robotics platforms' sensors and computer vision capabilities have made them viable platforms for K-12 AI programming. Engineers have also designed learning tools geared toward younger learners. Any-Cubes (Scheidt and Pulver, 2019), a prototype toy made of wooden blocks, a camera, and a microcontroller was designed to allow children as young as six explore image classification. In addition to the variety of learning tools aimed at enabling novice learners to create AI artifacts, an array of AI curricula designed for K-12 students has been developed in recent years.

AI Curricula and Approaches

As many K-12 educators have had little exposure to AI or experience teaching the subject, it is crucial that learning materials are approachable and relatable not only to students but also to teachers of all background levels (Vazhayil et al., 2019). Further, much of the existing AI curricula has been developed for college-level and other adult audiences; however, these learning materials are not always suitable for K-12 learning environments. Rigid class schedules in schools, broad learner developmental variability, and a wide range of unique factors require AI curriculum developers to consider strategies to maximize student engagement and provide the necessary scaffolding to help K-12 learners succeed (Zhou et al., 2020). Curriculum developers have set out to create curricula and other instructional materials to meet this need; however, few empirical studies have conducted rigorous evaluation of this budding AI curricula to determine their effectiveness.

Lin and Van Brummelen (2021) empowered teachers to customize AI curriculum to their students' interests. They engaged 15 5-12th grade teachers in co-design workshops, helping them to recognize opportune topics in the core subjects they taught

where AI could be embedded. The teachers and researchers then worked together to develop learning activities to help learners make connections between the core content topics and AI concepts. Lin and Van Brummelen suggested that in future curriculum co-design efforts, further scaffolding would better allow K-12 teachers to embed ethics and data concepts into their lessons. Additionally, the authors stressed the importance of curriculum developers providing ample support for student assessment and a focus on maximizing student engagement.

Other curriculum development initiatives have employed a variety of approaches to engage K-12 students with AI learning opportunities. In one such project, an AI researcher and classroom teachers co-designed an interdisciplinary curriculum aimed at grades K-6. The learning materials allowed learners to program Lego Mindstorms robots, write AI programs, and conduct experiments to deepen their understanding of basic AI concepts (Heinze et al., 2010). Ng, Lou, et al. (2022) designed a series of workshops designed to teach elementary-aged learners about AI through digital story writing. In another effort, Lee and Perret (2022) wrote a curriculum to introduce high school students to AI through data science. The lessons were co-designed with undergraduate students and geared toward learners with no background in AI.

In an effort to position learners at the heart of the learning experience, Toivonen et al. (2020) created a curriculum to engage middle school students with real-world ML. After the introductory lessons, students brainstormed problems that could be solved with the technology and collaborated with a group of their peers to build a prototype solution. To pique the interest of 10- to 14-year-old learners, Henry et al. (2021) developed a curriculum that incorporated ML learning opportunities into a role-playing

game. Students were introduced to ML concepts while engaging in game play. The curriculum was tested with 60 teachers, 70 middle school students, and 12 elementary students. Through pre- and post-tests as well as participant interviews the researchers found that learners deepened their understanding of the influence that training datasets have on AI systems' intelligence and that the technology's abilities are dictated by their underlying programs. Similarly, Lee et al. (2020) employed game-based learning to introduce AI concepts in upper elementary learning environments. They built a collaborative game-based platform designed to incorporate AI learning opportunities into the life science field through a problem-based approach. Students learned and applied image recognition, ML, and automated decision making methods through block-based programming.

Much of the vital work on K-12 curriculum development has been limited in scope and targeted at isolated components of the AI field. Rizvi et al. (2023) conducted a systematic literature review of 28 studies published from 2019 to 2022 that presented data on AI learning outcomes. Analysis of the research revealed that many efforts zeroed in on specific AI topics, leading the authors to suggest future research on evaluation frameworks that could inform curriculum developers' efforts to write well-rounded, comprehensive learning materials.

In addition to developing students' understanding of AI concepts, scholars have also emphasized the importance of embedding ethics instruction in the AI education.

Ethics and Equity

Due to the influence AI technology has on society and the risks it poses to fairness, equity, privacy, security, job stability, economic equality, and accurate and

truthful news reporting, there is increasing recognition of the value of having K-12 students reflect on the ethical issues posed by AI. In an article introducing activities to help students reflect on ethical decision-making using scenarios involving self-driving cars, Furey and Martin (2019) made a case for a modular method to AI ethics instruction. They argued that modular ethics curricula allow for the lessons to be adopted into existing AI courses in order to fill gaps in ethics instruction without requiring a complete rewrite of the course material. However, Borenstein and Howard (2021) contended that ethics is far too critical to be treated as something that is plugged in an existing, disconnected curriculum. Instead, they argued that it should be weaved in at every point that students learn about AI. The authors made a case that learners should experience lessons stressing ethics using different approaches throughout their coursework, including in technical classes.

Akgun and Greenhow (2021) suggested treating students' lived cultural experiences and perspectives as assets in AI ethics instruction. They encouraged curriculum to incorporate more culturally relevant and responsive pedagogies that create inclusive learning environments, celebrate students' backgrounds, and nurture learners' social awareness while reflecting on key ethical issues (Ladson-Billings, 1995). Akgun and Greenhow also emphasized that improvement in AI ethics instruction hinges on adequate teacher professional development (PD). Pre- and inservice teachers need opportunities to interact with potential curriculum resources, consider teaching best practices, and engage in communities of practice with peer teachers in order to build their capacity.

Blakeley Payne (2019), assisted by Cynthia Breazeal's MIT Media Lab Personal Robots Group, wrote an influential, freely available curriculum geared toward middle school students titled AI + Ethics. The curriculum's eight learning activities are largely unplugged lessons aimed at illustrating AI's positive and negative impacts on society. Goals of the curriculum included helping students comprehend how AI systems work, that they are not neutral, that not all audiences are impacted the same way by the technology, and to empower learners to seek equitable outcomes from the systems.

Zhang, Lee, et al. (2022) crafted the Developing AI Literacy (DAILy) workshop, a 30-hour middle school curriculum focused on three main topics: AI technical concepts and processes, ethical and societal implications of the technology, and careers in the coming AI age. The curriculum was designed to broaden AI participation by focusing on middle school students, an age the authors identified as pivotal for developing AI literacy. Due to the COVID pandemic, the workshop was presented virtually during a summer STEM program to 25 participants, the majority of whom had no previous CS/AI instruction. The sample was made up predominantly of students from groups traditionally underrepresented in CS and STEM fields. Results of an AI concept inventory suggested significant increases in students' comprehension of AI concepts, technical skills, and ability to recognize AI. Additionally, students demonstrated greater ability to spot and mitigate bias in ML models and also speak to the technology's impacts on society (Lee et al., 2021). The group's emphasis on embedding ethics into AI instruction has been mirrored by a range of other efforts.

In another study, Ali et al. (2019) piloted their AI+Ethics curriculum with 225 middle school students who had no previous AI or CS experience. The series of three

45-minute constructionist-styled lessons focused not only on AI concepts but also algorithmic bias. The researchers' approach demonstrated potential for embedding opportunities for students to reflect on the ethics of AI while engaging in learning about the field.

Pedagogy and Instruction

In addition to exploring what K-12 students should be taught, researchers have also examined which instructional approaches are most effective. In a systematic review of 32 AI education empirical studies published from 2010 to 2022, Yue et al. (2022) encountered a wide variety of pedagogies and instructional techniques employed in K-12 learning environments. Most commonly found in the research were direct instruction, interactive learning, collaborative learning, participatory learning, game-based learning, project-based learning, and design-oriented learning.

Prior research has suggested that K-12 teachers' instruction should be informed by their students' conceptions and background knowledge (Qian & Lehman, 2017; Sadler et al., 2013). Within AI education, research on learners' conceptions is nascent but initial studies suggest its implications for pedagogy. Better understanding of learners' conceptualization of the technology can inform teaching practices to supplement students' preconceptions and counteract their misconceptions. For example, Sulmont et al. (2019) interviewed 10 university professors from Canada, the U.S., and Europe who were experienced in teaching ML to novice students. They reported that the most common preconceptions among their college students originated from popular media and while students generally felt that the technology was powerful, will be increasingly important going forward, and could lead to future career

opportunities, many students did not think they were capable of harnessing ML because they did not possess strong enough CS and math skills. Long & Magerko (2020) reviewed 150 AI articles and found that learners commonly have misconceptions about the technology's limitations, wrongly assume that AI systems have much more intelligence than they actually do, and also have limited conceptions of how the systems function due to a lack of transparency. Meanwhile, AI preconceptions appear persistent. During Zhang, Lee, et al.'s (2022) implementation of an AI curriculum with middle school students, the researchers reported that of those learners who entered the learning experience with AI misconceptions, 37.5% were able to overcome them and develop a complex understanding of the content. However, the incorrect preconception persisted through the conclusion of the workshop for 25% of participants, suggesting both the promise and challenge of overcoming preconceptions.

As AI education begins to mature, researchers have begun to turn to established pedagogical frameworks, such as TPACK (Koehler & Mishra, 2009), to explore what teacher knowledge and which approaches to instruction are most advantageous. In their systematic review, Yue et al. (2022) found an increase in the number of studies scrutinizing teachers' technological pedagogical knowledge (TPK), specifically how they apply the affordances that new AI teaching tools and technologies present to the learning environment. Kim et al. (2021) used a collective case study approach to analyze the technological knowledge (TK), pedagogical knowledge (PK), and content knowledge (CK) required to implement commonly used AI resources and tools in U.S., Chinese, Indian, Australian, and South Korean classrooms. Their analysis informed their proposed set of 11 K-12 AI teacher competencies found in Table 1.

Table 1*K-12 AI Teacher Competencies*

TPACK	Required competency of teachers for K-12 AI education
PK	1. Facilitate project-based learning 2. Organize essential concepts and principles into play and games
PCK	3. Classroom management and problem-solving based activity construction on AI technologies and AI social issue awareness
CK	4. Fundamentals of AI (problem-solving, inference, learning, recognition) 5. Computer Science (programming, algorithms) 6. Applied Mathematics (probability, statistics, calculus) 7. AI Ethics
TK	8. Using ICT tools and educational software 9. Construct programming environment
TCK	10. Using web or API-based online education platforms for AI project education
TPK	11. Provide feedback and encourage peer-reviewing of AI project outcomes shared in open online education platform

Source: Kim et al., 2021

An important aspect of K-12 pedagogy is student assessment that can serve to inform teaching practices. Scholars have begun to close the research gap in the K-12 AI assessment space but much work is left to be done. Ng et al. (2023) developed a validated instrument aimed at measuring K-12 students' AI literacy. Their 25-item questionnaire used a 5-point Likert scale to gauge different components of students' AI literacy development and perception of their abilities. The researchers piloted the instrument at two schools in Hong Kong with 363 participants aged 12 to 17 and

conducted confirmatory factor analysis. Their model resulted in a good fit, with Cronbach's Alpha levels between 0.58 to 0.91, using the following factors: intrinsic motivation, self-efficacy, behavioral intention, behavioral engagement, know and understand AI, and use and apply AI.

Sound pedagogy is also responsive to student needs and outcomes. While examinations of differences in learning outcomes based on gender have been rare in the AI education space, Norouzi et al. (2020) reported a concerning observation. The authors held an 80-hour, one-month summer course for 19 high school students on ML and natural language processing. The instructors randomly assigned participants peer partners throughout the course, had students engage in icebreaker activities regularly, and frequently encouraged participants to have discussions with their partners and the whole class in order to encourage cooperative learning. However, when students shifted to mixed-gender group work, the young women generally opted for less technical tasks on projects and presentations while the young men in their groups largely took on the technical roles. While only one finding, this alarming report is a vivid reminder of the importance of not just creating opportunities for students from traditionally underrepresented groups to participate in AI learning, but to prioritize instructional approaches that build learners' self-efficacy and leverage their interests, strengths, and lived experiences.

Underrepresented Groups

It has been well-documented that White, Asian, and/or male learners have greater access to and higher participation rates in CS/AI education than their non-male and/or minority peers (Zhang, Lee, et al., 2022). Meanwhile, it is crucial that all learners,

especially those from traditionally underrepresented groups, receive AI education in order to improve the inclusivity of AI design and accessibility of the tools that have become key in our digital age (Ng et al., 2021a). Researchers and other organizations have begun to implement initiatives to boost inclusivity in AI education.

AI4ALL (2020), a non-profit organization based in the U.S. seeks to provide learning opportunities to students from traditionally excluded groups. They offer a free interdisciplinary curriculum aimed at high school students that is aligned to science, language arts, and CS standards. The ten-lesson curriculum was written with an eye toward accessibility to all students and does not require students to have extensive math and CS backgrounds. Additionally, the organization partners with 16 postsecondary U.S. universities to provide programs to students of diverse backgrounds and identities in an effort to broaden participation in AI so that future technologies better reflect the entire society and its needs. Eguchi (2021), recognized the need to make access to AI learning opportunities widely available in order to prevent a deepening of the digital divide between developed and developing nations. The researcher developed Cogbots, an inexpensive open source robotics platform for students to learn AI. To keep costs as minimal as possible, robots can be assembled using recycled smartphones and 3D printed or cardboard components, in addition to inexpensive circuitry.

Everson et al. (2022) sought to create an opportunity for students to be critical computer scientists. The result was a six-week course for high schoolers titled “Creatively Coding a Better Future” that was co-constructed with a group of 14 racially, ethnically, socioeconomically, and gender diverse students from the state of Washington. Using a culturally responsive pedagogy, the instructors allowed for

maximum student agency on projects and students responded by creatively weaving critical issues into their work for each of the CS and AI topics covered. The researchers emphasized that when approaching a potentially sensitive topic such as ethics in artificial intelligence, it is necessary to establish a classroom climate in which students feel comfortable discussing topics related to underrepresented groups.

K-12 AI and ML Curricular Frameworks

The AI field is evolving quickly, making it challenging for schools to respond with appropriate learning opportunities for their students. As AI topics have yet to make their way into U.S. K-12 content standards in a meaningful way, educators have lacked an authoritative source to turn to as they consider what to teach about AI. In an attempt to fill this void, a variety of K-12 AI and ML curriculum frameworks have recently been developed, including AI Literacy: Competencies and Design Considerations, the AI4K12: K-12 AI Guidelines, and the Machine Learning Education Framework for Transforming ML Consumers to be ML Contributors.

AI Literacy: Competencies and Design Considerations

As multimedia became more commonplace in the 1990s, experts began to suggest the need for a new skill set which would allow the future workforce to take full advantage of the digital tools and information sources rapidly coming available. Gilster popularized the term “digital literacy” in 1997 and galvanized an effort to provide students opportunities to learn the skills and competencies to help them succeed in the digital world. Gilster’s view of digital literacy centered on helping learners nurture new ways of thinking rather than memorizing technical procedures. Subsequently, as the world has quickly adopted a variety of digital tools, digital literacy has become regarded

as critical to individuals' career success, schools' suitability for the times, and nations' competitiveness with peer countries (Bawden, 2008). Many are pressing for a similar emphasis on AI literacy today. Kandlhofer et al. suggested that, "Looking towards the near future, jobs will largely be related to AI. In this context literacy in AI and computer science will become as important as classic literacy (reading/writing)" (2016, p. 1).

Long and Magerko's (2020) seminal work argued that despite the increased use and broader acceptance of AI, the greater citizenry does not have a firm understanding of how the systems function. Extensive adoption of AI systems could have myriad implications on society, and these impacts could become more pronounced if the general public, including policymakers, have a limited understanding of the technology. As a result, individuals would become less competitive in the workplace and greater society could become increasingly unjust. To combat these potentially dangerous outcomes, the authors called for measures to boost AI literacy education in order to deepen society's level of understanding of AI. In an attempt to jump start these efforts, the researchers' first work was to establish what everyone should know about AI.

Due to the dearth of peer-reviewed AI education research, Long and Magerko's scoping study (2020) broadened to consider syllabi from AI courses, AI textbooks, polls of public perception, and research examining common perceptions and misconceptions of AI. Their analysis led to themes that were distilled to a working definition of AI literacy in addition to a framework of learner competencies and design considerations to inform the design of AI literacy learning experiences. Akin to terms digital literacy, computational literacy, scientific literacy, and data literacy used to denote what individuals comprehend and the actions they are able to perform in a certain context,

Long and Magerko define *AI literacy* as “a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace” (2020, p. 2). The authors also proposed 17 learner competencies (see Table 2) cataloging the knowledge, skills, and abilities they feel every learner, especially non-technical learners, should be able to demonstrate following AI instruction.

Table 2

AI Literacy Competencies

Competency	Learner Outcome
1. Recognizing AI	Distinguish between technological artifacts that use and do not use AI.
2. Understanding Intelligence	Critically analyze and discuss features that make an entity “intelligent”. Discuss differences between human, animal, and machine intelligence.
3. Interdisciplinarity	Recognize that there are many ways to think about and develop “intelligent” machines. Identify a variety of technologies that use AI, including technology spanning cognitive systems, robotics and ML.
4. General vs narrow AI	Distinguish between general and narrow AI.
5. AI strengths and weaknesses	Identify problem types that AI does/does not excel at. Determine when it is appropriate to use AI and when to leverage human skills.
6. Imagine future AI	Imagine possible future applications of AI and consider the effects of such applications on the world.
7. Representations	Understand what a knowledge representation is and describe some examples of knowledge representations.
8. Decision-making	Recognize and describe examples of how computers reason and make decisions.
9. ML steps	Understand the steps involved in machine learning and the practices and challenges that each step entails.

Table 2 (cont'd)

10. Human role in AI	Recognize that humans play an important role in programming, choosing models, and fine-tuning AI systems.
11. Data literacy	Understand basic data literacy concepts.
12. Learning from data	Recognize that computers often learn from data (including one's own data).
13. Critically interpreting data	Understand that data requires interpretation. Describe how the training examples provided in an initial dataset can affect the results of an algorithm.
14. Action and reaction	Understand that some AI systems have the ability to physically act on the world. This action can be directed by higher-level reasoning (e.g., walking along a planned path) or reactive impulses (e.g., jumping backwards to avoid a sensed obstacle).
15. Sensors	Understand what sensors are and that computers perceive the world using sensors. Identify sensors on a variety of devices. Recognize that different sensors support different types of representation and reasoning about the world.
16. Ethics	Identify and describe different perspectives on the key ethical issues surrounding AI (i.e., privacy, employment, misinformation, the singularity, ethical decision making, diversity, bias, transparency, accountability).
17. Programmability	Understand that agents are programmable.

Source: Long and Magerko, 2020

In addition to the learner competencies, Long and Magerko's AI literacy framework also includes 15 design considerations (see Table 3) geared toward providing guidance to developers and educators as they construct learning experiences, support materials, and tools for AI learners.

Table 3*AI Literacy Design Considerations*

Design Consideration	Consideration Details
1. Explainability	Include graphical visualizations, simulations, explanations of agents' decision-making processes, or interactive demonstrations in order to aid learners' understanding of AI.
2. Embodied interactions	Design interventions in which individuals can act as or follow the agent, as a way of making sense of the agent's reasoning process. This may involve embodied simulations of algorithms and/or hands-on physical experimentation with AI technology.
3. Contextualizing data	Encourage learners to investigate who created the dataset, how the data was collected, and what the limitations of the dataset are. This may involve choosing datasets that are relevant to learners' lives, are low-dimensional and are 'messy' (i.e. not cleaned or neatly categorizable).
4. Promote transparency	Promote transparency in all aspects of AI design (i.e. eliminating black-box functionality, sharing creator intentions and funding/data sources, etc.).
5. Unveil gradually	To prevent cognitive overload, give users the option to inspect and learn about different system components; explain only a few components at a time; or introduce scaffolding that fades as the user learns more about the system's operations.
6. Opportunities to program	Provide ways for individuals to program and/or teach AI agents. Keep coding prerequisites to a minimum by focusing on visual/auditory elements and/or incorporating strategies like Parsons problems and fill-in-the-blank code.
7. Milestones	Consider how perceptions of AI are affected by developmental milestones (e.g. theory of mind development), age, and prior experience with technology.
8. Critical Thinking	Encourage learners to be critical consumers of AI technologies by questioning the intelligence and trustworthiness of AI applications.

Table 3 (cont'd)

9. Identities, values and backgrounds	Consider how learners' identities, values, and backgrounds affect their interest in and preconceptions of AI. Learning interventions that incorporate personal identity or cultural values may encourage their interest and motivation.
10. Support for parents	When designing for families, help parents scaffold their children's AI learning experiences.
11. Social interaction	Design AI learning experiences that foster social interaction and collaboration.
12. Leverage learners' interests	Exploit current issues, everyday experiences, or common pastimes like games or music when designing AI literacy interventions.
13. Acknowledge preconceptions	Allow for the fact that learners may have politicized or sensationalized preconceptions of AI from popular media, and consider how to respect, address, and expand on these ideas in learning interventions.
14. New perspectives	Introduce perspectives that are not as well-represented in popular media (e.g. less publicized AI subfields, balanced discussions on the dangers and benefits of AI).
15. Low barrier to entry	Consider how to communicate AI concepts to learners who do not have extensive backgrounds in mathematics or computer science (e.g. by reducing the prerequisite knowledge/skills, relating AI to prior knowledge, and addressing learners' insecurities about their ability).

Source: Long and Magerko, 2020

AI4K12

Non-profit advocacy organizations, government agencies, and industry partners have galvanized their efforts to bring opportunities to learn CS to all U.S. K-12 students. These efforts have also helped to institutionalize K-12 learner outcomes through the K-12 CS Framework (K-12 Computer Science Framework Steering Committee, 2016), CSTA K-12 Standards (Computer Science Teachers Association, 2017), AP CS Principles Curriculum Framework (College Board, 2020), and other guidelines. As an

outgrowth of this work, a variety of curricula, supporting resources, and educational tools have been developed to support implementation of CS instruction in the K-12 classroom. While these efforts have had a dramatically positive impact on the number of K-12 learners receiving CS education, the benefit has not spilled over into K-12 AI education. In an effort to replicate the successful standardization of what learners should know about CS, the Associate of Advancement of Artificial Intelligence (AAAI) and CSTA partnered to form the AI for K-12 Working Group (AI4K12). The working group is made up of professors, researchers, industry experts, and K-12 teachers and intends to map out the essential AI topics, concepts, and issues that should be taught to all students (Touretzky et al., 2019).

The AI4K12 Initiative developed a draft framework based on five big ideas (Wiggins & McTighe, 2005) about AI that all students should comprehend (see Table 4). The group then broke each big idea into learning concepts and skills before mapping them to learning objectives and essential understandings tied to K-2, 3-5, 6-8, and 9-12 grade bands. The group is in the process of receiving and incorporating feedback from the general public on the draft guidelines. AI4K12 envisions that their framework of learning pathways will serve to inform standards writing and curriculum development efforts to incorporate age-appropriate AI concepts and skill development (Gardner-McCune et al., 2019).

Table 4

AI4K12's Five Big Ideas

Big Idea	Idea Details
1. Perception	Computers perceive the world using sensors
2. Representation and Reasoning	Agents maintain representations of the world and use them for reasoning
3. Learning	Computers can learn from data
4. Natural Interaction	Intelligent agents require many types of knowledge to interact naturally with humans
5. Societal Impact	AI can impact society in both positive and negative ways

Source: Touretzky et al., 2019

The Machine Learning Education Framework

Lao established the Machine Learning Education Framework in a 2020 study. Where the AI Literacy and AI4K12 frameworks focus on laying out the concepts K-12 learners should understand and what skills they should have in the greater AI field, the aim of the ML Education framework was to transform “ML consumers to ML-engaged citizens, and then onwards to tinkerers and eventually researchers or engineers” (Lao, 2020, p. 149). While the author did not detail the literature review process that informed the framework, the ML Education framework was influenced by self-efficacy (Bandura, 1977), constructionism (Papert, 1980), and computational action (Tissenbaum et al., 2019) theories. The framework is concentrated specifically on the AI subfield of ML and extends past the K-12 setting and into higher education in its consideration of some technical ML concepts. The ML Education framework was intended to inform curriculum development efforts, appropriate evaluation of ML interventions, and the delivery of equitable learner supports. Lao contended that for learners to succeed and progress

from consumer to tinkerer to engaged citizen in the field, ML learning experiences must provide learners opportunities to cultivate the appropriate knowledge, skills, and attitudes (see Table 5). The framework’s four knowledge items are General ML, ML Methods, Bias in ML Systems, and Societal Implications of ML; the six skill items are ML Problem Scoping, ML Project Planning, Creating ML Artifacts, Analysis of ML Design Intentions and Results, ML advocacy, and Independent Out-of-Class Learning; and the four attitudinal items are Interest, Identity and Community, Self-Efficacy, and Persistence (Lao, 2020).

Table 5

The Machine Learning Education Framework for Transforming ML Consumers to be ML Contributors

Item	Learner Outcome
Knowledge	
1. General ML Knowledge	Know what machine learning is (and is not). Understand the entire pipeline of the creation of ML systems.
2. Knowledge of ML Methods	Identify when to use a range of ML methods across the breadth of the field (e.g., k-nearest neighbors, CARTs or decision trees, neural networks, ensemble methods). Understand how different methods work.
3. Bias in ML Systems	Understand that systems can be biased, and the different levels and ways in which bias can be introduced.
4. Societal Implications of AI	Understand that ML systems can have widespread positive and negative impacts. Consider the ethical, cultural and social implications of what they do.
Skills	

Table 5 (cont'd)

1. ML Problem Scoping	Determine which problems can and should be solved by ML.
2. ML Project Planning	Plan a solution which is sensitive to both technical and contextual considerations.
3. Creating ML Artifacts	Use tools to create appropriate artifacts.
4. Analysis of ML Design Interactions and Results	Describe the explicit and implicit design intentions of an ML system. Critically analyze the intentions against how the system can and should be used.
5. ML Advocacy	Critically discuss ML policies, products and education.
6. Independent Out-of-class Learning	Students seek learning experiences outside the classroom.

Attitudes

1. Interest	Students are engaged and motivated to study the topic.
2. Identity and Community	Students contribute to and learn from a community of peers and/or broader online communities who are interested in ML.
3. Self-efficacy	Students are empowered to build new, meaningful things.
4. Persistence	Students continue and expand their engagement with ML.

Source: Lao, 2020

Lao was mindful that not all schools would have the resources necessary for learners to create the technical artifacts necessary to address all items in the framework. To account for these under-resourced situations, the framework requires three of the knowledge items and three skill items as a minimum to help learners become ML-engaged citizens, with hope that they can be exposed to the technical knowledge and skills in the future. These minimally required items are K1. General ML Knowledge, K3. Bias in ML Systems, K4. Societal Implications of ML, S1. ML Problem

Scoping, S4. Analysis of ML Design Intentions and Results, and S5. ML Advocacy (Lao, 2020).

Teacher Perceptions and Conceptions

Teachers are indispensable in any widespread effort to provide AI learning opportunities in K-12 schools. However, in their analysis of 49 research articles, Ng, Lee, et al. (2022) found that K-12 teachers commonly expressed low confidence in their grasp of AI. Given teachers' integral role in AI education, their perceptions and conceptions of the technology could be critical, yet a small number of documented studies have deeply examined K-12 teachers' interpretations and views.

Vazhayil et al. (2019) examined middle and high school teachers' conceptions about teaching AI as a part of teacher PD in India. The authors found that the majority of teachers did not recognize the job opportunities that AI learning could create for their students. Teachers' primary concerns focused on the challenge of communicating the importance of the field to school administrators and the lack of access to AI instructional tools. Participants reported that access issues were largely due to the high level of internet filtering on school computers.

Lindner and Berges (2020) used semi-structured interviews to explore how 23 German teachers' who taught a variety of subjects including CS conceived of AI. Only two participants felt confident in their ability to teach AI while nine reported that they had independently sought out information to learn more about AI. Teachers frequently felt that AI imitates humans' thinking and exhibits creativity and problem-solving skills. Teachers often identified AI technologies to be algorithms and programs that also have the ability to self-learn and grow abilities beyond their coded instructions through

identifying patterns in data and trial and error. The researchers concluded that based on educators' vague and general conceptions of AI, they would need to develop deeper content knowledge through teacher PD in order to successfully teach the topic. Despite the importance of teachers learning about AI, only 40% of participants reported actively learning about the field, while the majority of the teachers' knowledge came from popular media and scientific reports.

Chiu and Chai (2020) interviewed 24 Hong Kong computing teachers to examine their perceptions of K-12 AI education. Participants taught grades 7-12 and half had previous experience teaching AI while the other half had no experience teaching the subject. The researchers found that while the teachers generally felt it was important for students to learn about AI, they considered themselves underqualified and that they did not have an adequate level of knowledge to teach the content. Teachers also did not feel they had enough time in their rigid courses to add AI content to their curriculum, nor did they believe that they had access to the pertinent tools necessary for teaching the subject.

To better understand teachers' preconceptions of teaching ML, Sanusi et al. (2022) conducted a phenomenological study with 12 inservice African teachers. The purposively sampled participants had all earned a computing degree and came to the study with experience teaching 10-12th grade CS courses in geographically diverse locations in Nigeria, Ghana, Tanzania, Kenya, South Africa, and Ghana. Sanusi et al.'s analysis of the teachers' responses during semi-structured interviews led to the following five themes: supporting student technical knowledge, having knowledge of the concept, focusing on PD practices, contextualizing teaching resources and tools, and

sustainability for development goals. Teachers most frequently spoke of the need for PD since even though they were familiar with ML they did not feel they had the necessary technical knowledge to teach the subject. Teachers also stressed the importance of ensuring that ML learning opportunities and materials are appropriately tailored to the students being taught.

Ottenbreit-Leftwich et al. (2022) explored elementary school teachers' conceptions of AI and their approach to writing individualized curriculum based on students' comprehension. The researchers interviewed three 4th and 5th grade teachers and asked them to define AI, whether they felt it should be taught, and how to best blend ethics into the instruction. Teachers pointed out that they regularly encounter and recognize AI in the form of personal assistants such as Siri and Alexa but that it is more challenging to detect technologies that often incorporate AI less overtly. Participants also explained that while AI is programmed by humans, the systems are able to make decisions beyond the instructions they are explicitly given, often through analyzing data. They generally felt that AI should be taught in part because new state standards require it but also to help their students see how the technology already surrounds them. Teachers underscored that as AI's influence grows in the future they want to be sure that their students are prepared for success. In addition, participants emphasized the importance of ethics instruction and felt that data privacy, self-driving cars, and character education would be fruitful methods to address AI ethics.

Yau et al. (2023) analyzed teachers' conceptions of teaching AI through a phenomenological study. After selecting 28 inservice teachers in 17 schools representative of Hong Kong, the authors co-designed a curriculum with the teachers.

The teachers implemented the curriculum with 7th through 9th grade students and participated in interviews exploring their experience and views. The study categorized teachers' conceptions of the importance of AI instruction into the following six themes: technology bridging, knowledge delivery, interest stimulation, ethics establishment, capability cultivation, and intellectual development. The most commonly discussed themes were technology bridging, or increasing students' awareness of the AI technology they encounter in their lives, and knowledge delivery, or teachers' role in shaping the content into learning opportunities appropriate for their students.

Ayanwale et al. (2022) explored potential relationships between teachers' conceptions of AI and their intention to teach the subject. Their study surveyed 385 Nigerian K-12 teachers and used structural equation modeling to analyze participants' responses to the online survey. Thirty-three items and the following eight variables were applied: AI anxiety, perceived usefulness of AI, AI for social good, attitude towards using AI, perceived confidence in teaching AI, relevance of AI, AI readiness, and behavioral intention to teach. The researchers found that teachers' perceived confidence in teaching AI is a meaningful predictor of their intention to teach the subject, while teachers' belief in the relevance of teaching AI predicts their readiness to teach the subject.

Holes in Research Base

As appeals for AI education in grades K-12 increase, it is imperative that teachers are prepared and equipped to lead this instruction. Additionally, for K-12 AI education to happen effectively, decision makers must build programs of study thoughtfully, drawing on empirical research. The existing lines of research offer valuable

initial insights into the field; however, few exploratory, descriptive studies have examined 9-12th grade AI education in the U.S. CS courses. The bulk of extant AI education research focuses on postsecondary settings leaving holes in the research base at the K-12 level. A growing number of studies have examined K-12 teacher conceptions of AI but their samples are typically drawn from the general K-12 teacher population. It is crucial that CS teacher conceptions of AI are analyzed for two primary reasons. First, AI is a sub-field of CS. Second, today's CS students may become tomorrow's AI developers, and it is imperative that they are well grounded in the technology's equity and bias dilemmas.

The first step to build this research base is to understand how CS teachers' conceive of AI and include it in their CS instruction. However, largely missing from the literature on AI education are studies examining 9-12th grade CS teachers' conceptions and understanding of the field, their perceptions of the domain and its role in secondary education, their confidence in teaching AI concepts and practices, and the challenges and roadblocks they face in teaching the subject. The goal of this study was to narrow this gap in the research by exploring teachers' conceptions of AI education in high school CS courses through the following research questions:

RQ1: How do 9-12th grade CS educators conceptualize AI and ML?

RQ2: How do 9-12th grade CS educators conceptualize the role of AI in their CS instruction?

RQ3: What, if any, instructional curricula and pedagogies do 9-12th grade CS educators use to bring AI into their CS instruction?

CHAPTER 3: METHODS

Positionality Statement

Kezar (2022) suggested that individuals do not have one exclusive identity but a variety of overlapping identities and such is the case with me. I am a white male who attended public schools in the U.S. throughout grades K-12 and college. Prior to embarking on this research, I taught CS full-time in grades 6-12 at public schools in the U.S. for 19 years. I also provided PD to inservice teachers who were formally prepared to teach other subjects (e.g., science, math, history, language arts) but shifted to teach CS. Knowing that researchers and study participants each heavily influence the research process (England, 1994), I understand that my experience in K-12 CS education could jeopardize participants' ability and comfortability to express their conceptions openly. Bourke (2014) urged that it is crucial for researchers to recognize and acknowledge their identities and accompanying biases. This recognition alone can serve to provide insights into their positionality and the steps necessary to develop a plan to successfully and objectively conduct the research. Similarly, Malterud (2001) underscored that identifying and acknowledging these types of experiences and potential preconceptions are key to a reflexive approach to successfully design and conduct impartial research. While I was formerly a member of the population being studied (9-12th grade U.S. CS teachers), I am no longer. However, I felt it important to acknowledge this background to the study's participants.

In the later stages of my teaching career, I felt a strong sense that my students needed to be prepared to enter an increasingly AI-saturated world but was unsure of the best way to provide this preparation. The frustration of this experience during my time in

the CS classroom was a primary motivation for conducting this research. These personal questions certainly informed my inquiry in this study, but I made a deliberate effort to develop well defined, structured research methods to prevent my background from contaminating the study in ways that might lead the research to reinforce my preexisting notions; those methods follow.

Recruiting and Participants

The study sought teachers who were currently teaching 9-12th grade CS at in-person or formal virtual schools based in the United States. High school teachers were the aim of the study since AI concepts are currently only found at the 11-12th grade level of the latest version of the CSTA (2017) student standards. While many K-8 CS teachers undoubtedly introduce AI concepts to younger students, bounding the sample to 9-12th grade teachers also allowed the study to have greater focus and consistency.

A request for participation was shared with CSTA's teacher network through their international online forum and local teacher community chapters. CSTA was chosen as the primary recruitment avenue because it is a U.S.-based, teacher-led non-profit organization geared toward K-12 CS educators. The recruiting messaging (see Appendix A) communicated the purpose of the study, expectations of participants, and that participants would be compensated with a \$100 gift card for their participation in the study. The message also detailed the criteria for participation, provided a link to a background survey for teachers to volunteer, and shared my contact information so potential participants could seek more information on the study.

Fifty-three teachers expressed their interest in participating in the study by completing the participant screening survey (see Appendix B). The questionnaire first asked potential participants the following information about their teaching background and current position: how long they had been teaching CS content, how they prepared to teach CS content, what CS courses they taught, the grades of their CS students, what non-CS courses they taught, whether and how they incorporated CS content into the non-CS courses, and the grades of their non-CS students. The survey then collected each volunteer's name, email address, preferred pronouns, race / ethnicity, school student enrollment, school setting, and school's Title I status. Responses were examined to ensure that potential participants met the sample criteria, and three respondents were removed from consideration as their answers suggested that they did not teach CS content to 9-12th grade students at a U.S. school. The remaining 50 potential teachers were emailed an invitation to schedule an interview via a web-based calendaring tool. The original participant recruitment strategy was not to achieve a predetermined sample size, but rather the intent was to collect data from teachers until data saturation was reached and no new themes continued to emerge from the analysis of the interviews (Lincoln & Guba, 1985). After interviewing 23 teachers, I conducted an initial data analysis. No novel, substantive points arose during the coding of the final four transcripts of the dataset, so I determined that saturation had been reached.

Twelve participants indicated she/her/hers pronouns (52.2%), nine identified with he/him/his (39.1%), and two did not indicate preferred pronouns (8.7%). Sixteen teachers identified as White (69.6%), three as Hispanic or Latino (13.0%), one as Asian (4.4%), one as Black or African American (4.4%), and two as Multiethnic or Mixed

(8.7%). Three participants had 0-1 year experience teaching CS (13.0%), four had 2-4 years (17.4%), ten had 5-9 years (43.5%), and six teachers had more than 10 years experience teaching CS (26.1%). Four teachers taught at a school with enrollment between 0-249 students (17.4%), eight with enrollment between 250-1000 (34.8%), five with enrollment between 1000-2000 (21.7%), and six teachers taught in a school with an enrollment above 2000 students (26.0%). Eleven participants taught in school in an suburban setting (47.8%), eight in an urban school (34.8%), and four in a rural setting (17.4%). Fifteen teachers taught in non-Title I schools (65.2%), six in Title I (26.1%), and two were unsure whether their school was Title I (8.7%). See Table 6 for details on the 23 individual participants.

Table 6

Participant Characteristics (n=23)

Pseudonym	Pronouns	Race / Ethnicity	CS Teaching Experience	School Enrollment	School Setting	Title I School
Axel	he/him/his	Asian	5-9 years	250-1000	rural	no
Winonah	she/her/hers	White	10+ years	0-249	suburban	no
Elliott		Mixed	0-1 year	0-249	suburban	no
Liam	he/him/his	Multiethnic	10+ years	250-1000	suburban	unsure
Orin	he/him/his	White	2-4 years	1000-2000	urban	yes
Fatima	she/her/hers	White	10+ years	250-1000	rural	no
Damari	he/him/his	White	10+ years	2000+	suburban	no
Bella		White	5-9 years	250-1000	suburban	no
Violet	she/her/hers	White	10+ years	250-1000	suburban	no
Cora	she/her/hers	White	5-9 years	0-249	urban	No
Gemma	she/her/hers	Black or African American	5-9 years	2000+	urban	no

Table 6 (cont'd)

Keshawn	he/him/his	White	2-4 years	1000-2000	urban	yes
Henry	he/him/his	White	5-9 years	250-1000	rural	no
Penelope	she/her/hers	White	5-9 years	1000-2000	suburban	unsure
Nakeisha	she/her/hers	Hispanic or Latino	5-9 years	250-1000	rural	yes
Jose	he/him/his	Hispanic or Latino	2-4 years	1000-2000	suburban	no
Mia	she/her/hers	White	5-9 years	0-249	urban	no
Queen	she/her/hers	White	5-9 years	1000-2000	suburban	no
Itzel	she/her/hers	White	5-9 years	2000+	suburban	yes
Uri	he/him/his	White	0-1 year	2000+	urban	no
Sebastian	he/him/his	Hispanic or Latino	2-4 years	2000+	urban	yes
Tamika	she/her/hers	White	10+ years	250-1000	suburban	no
Ruby	she/her/hers	White	0-1 year	2000+	urban	yes

Potential interview participants were emailed the IRB-approved informed consent form (see Appendix C) that explained the purpose of the study, what participants would be asked to do, the potential benefits and risks of participating, privacy and confidentiality procedures, rights to opt out of or withdraw from the study, compensation, and my contact information. The form also explained the informed consent procedure in which participants indicate their voluntary agreement to participate in the study by beginning the video call interview. At the beginning of the interview meeting, participants were asked if they had any questions or concerns regarding the content of the consent form before the interview began.

Interview Protocol

A semi-structured interview protocol was used to examine teachers' conceptions of AI and ML, the role of AI in their CS instruction, and the curricula and pedagogies they use to bring AI into their CS instruction. Due to the scarcity of research on teachers' conceptions of AI and their role in CS instruction, new interview questions were developed specifically for the study. A pilot interview was conducted with a 9-12th grade CS teacher to seek feedback on the interview questions and process. The pilot interviewee felt the questions were understandable and appropriate and also provided suggestions for phrasing to improve question clarity and approaches to lower participants' apprehension about their content knowledge. While I strove for consistency across all interviews, the interview protocol was slightly adapted for individual circumstances, time constraints, and follow up questions. This emergent design approach allowed for interviews to adapt to details that were uncovered in the course of the interviews (Creswell & Poth, 2016). The interview protocol used for each interview is presented in Appendix D.

Each interview began with rapport building to help create a relaxed, accepting environment in an effort to make participants feel comfortable and that their insights were valued (Lichtman, 2006). This was especially important since the interviews were conducted remotely without physical proximity between the interviewer and interviewee (Roberts et al., 2021). I asked teachers what their favorite course was and to share about something exciting that had happened in one of their classes recently. To explore research question 1, participants were then asked how they would explain AI and then ML to someone who has never encountered the terms. Shifting gears to research

question 2, the interview protocol moved to questions about how they currently were or could envision teaching AI in their CS courses in terms of the topics they covered, takeaways for students, and challenges they encountered. Further questions examined the high school teachers' thoughts on AI being taught in K-12 schools including if teachers felt the topic belongs in schools and when and what should be taught to which students. Questions also sought teachers' insights on issues that are likely to arise with AI instruction as well as where they would prioritize AI among the myriad of CS topics at the K-12 level. The final portion of the interview delved into participants' experiences with AI pedagogical strategies and curricula in connection to RQ3. I then played the role of a teacher while participants played the role of students in a sample unplugged activity (see Appendix E). The lesson was modified from Lindner et al.'s (2019) activity designed to introduce ML classification models and decision trees where learners develop rules to classify images of cartoon animals. Adjustments were made to the lesson to make it more appropriate to complete in the online meeting setting with only one student rather than an entire class. Additionally, cow faces were used as opposed to the original monkey images due to the potential racial connotations associated with monkeys within the U.S. context. At the conclusion of the lesson, participants were asked how they felt the activity would be received by their students, how they felt unplugged activities would influence their students' understanding of AI, and what key considerations in the design and delivery of unplugged AI activities would benefit their students.

At the conclusion of the interview, teachers were asked whether they had other thoughts that were not yet discussed that they would like to share. This allowed

teachers space to add their thoughts on issues that they deemed critical that were not included in the formal interview protocol. Throughout the interviews probes such as “Can you tell me more about ...”, “Will you please clarify what you mean by ...”, “What else can you say about that ...”, and “Will you please provide more detail about ...” were used to elicit more details when participants provided terse responses or introduced novel ideas.

Data Collection Procedure

Interviews lasted between 55 and 90 minutes with an average interview lasting 65 minutes. All interviews were conducted remotely through virtual meeting software (e.g., Zoom, Google Meet). Hanna and Mwale (2017) have found remote video interviewing to be an efficient means to conduct interviews with distant participants without the expense and time required to travel for a face-to-face meeting. Additionally, a video interview allows both the participant and researcher access to visual cues and body language that strengthen rapport- and trust-building. The technologies’ video and audio recording features allow the researcher to focus on participants’ responses rather than having to divert attention to transcription or the risk of recording equipment failing. An added affordance of a video interview is that it can provide balance in the power relationship between the participant and researcher, as the participant may be less intimidated to speak to the researcher from a physical place of comfort, such as their own home or workplace, leading to a more open interview and insightful responses. The risk inherent in video interviews is that the technology becomes a hindrance to the conversation. Low-speed or intermittent internet connectivity, hardware malfunction, and software stability can all endanger the quality of data collected from the online interview

(Hanna & Mwale, 2017). Internet connectivity only posed a problem during one of the interviews in the study; however, I was able to reconnect with the participant after our video call was disconnected and complete the remainder of the interview.

I solely conducted all interviews during a five-week span in an attempt to maintain a consistent interview procedure. Participants were presented a choice between virtual meeting platforms so they could select the software platform that they were most comfortable with and to minimize technological challenges (Roberts et al., 2021). Participants were asked at the beginning of the interview whether they consented to being recorded, all consented, and an audio and video recording was made of each interview. This allowed me to focus on exploring participants' answers and asking probing follow-up questions during the interview rather than working exclusively to transcribe participants' responses real-time.

Data Analysis

Qualitative data analysis is not a solitary, one-time occurrence but an ongoing, cyclical practice that begins during data collection and continues throughout the course of the study (Miles & Huberman, 1994). A thematic analysis approach was chosen for the study's data analysis method due to its versatility and emphasis on seeking collective meaning across teachers' experiences throughout the data set. Braun and Clarke (2006), in their seminal work on thematic analysis, suggest researchers apply the following six-phase process: 1) Familiarizing yourself with the data, 2) Generating initial codes, 3) Searching for themes, 4) Reviewing potential themes, 5) Defining and naming themes, and 6) Producing the report. The following details each step in the

thematic analysis process and describes how it was applied to the current study's data analysis.

Phase 1: Familiarizing Yourself With the Data

The audio recordings were transcribed using Otter.ai, an automatic transcription tool, and participant names were replaced with pseudonyms. The text of the transcriptions was compared to the interview audio recordings to check their accuracy and the necessary corrections were made. The transcripts were then brought into the NVivo qualitative data analysis software to organize and manage the interview data and support the coding process. After reviewing the transcripts, I read through the dialogue of each complete interview to orient myself to the contents of the transcripts before beginning coding. Notes were kept of remarkable passages and ideas that might impact the future data analysis phases since coding is developed and refined throughout the analysis process (Braun & Clarke, 2006).

I wrote analytic memos throughout the coding process to document the thinking and choices I made and also to reflect on the patterns that were emerging. This led to memos that detailed connections between teachers' conceptions, possible new directions in the data analysis, thought-provoking participant quotes, and questions that I did not yet have answers for (Saldaña, 2021). My memo-writing process was guided by Charmaz's (2014) suggestion to write what occurred to me at the time and only later in the data analysis process seek to determine the relevance and importance of each of the captured reflections. This approach allowed me to focus on providing a running commentary of the thought process behind data analysis and the conclusions I drew,

working to strengthen the trustworthiness and credibility of the study's findings (Shenton, 2004).

Phase 2: Generating Initial Codes

As this was an exploratory study, it was important to make space for potential new concepts and theories to take shape. In order to do so, I leveraged an inductive approach to coding where the teachers' words gave birth to all of the codes rather than entering the coding process with predetermined codes that the participants' thoughts needed to conform to. Using an open coding technique for the first round of coding allowed for new concepts to emerge from the interviews and be discovered more organically. I used the In Vivo Coding technique, a first cycle grounded theory method, and wherever possible codes were created using participants' verbatim phrasing from the interviews. By generating codes directly from the actual language teachers used in question responses, my data analysis was able to stay closer to participants' experiences and original meanings (Strauss, 1987; Stringer & Aragón, 2020). Saldaña (2021) suggests that In Vivo Coding is well matched for studies that rely on interview data, especially when the focus is to center participants' voices and perspectives. As part of the In Vivo Coding process, I scrutinized interview transcripts with an eye toward systematically extracting understanding from each interview, organizing those interpretations across participating teachers' responses, and consolidating meaning throughout all of the interview data (Grbich, 2012).

After the initial codes were established, I reread each interview transcript and tagged phrases and passages associated with the newly created codes. During this cycle, when I encountered passages with additional noteworthy meaning that were not

captured by the initial set of codes, I created a new code based on the passage. This resulted in 138 total codes. Each code was defined with a phrase that described the code label in further detail. Since some of the codes grew out of the review of the interviews and a portion of them did not arise until near the end of the coding pass, I felt it was important to work through all interviews again in order to code them with the full complement of 138 codes. This ensured better coverage of the entire codebook to each and every transcript. The codes, their definitions, and the interview passages associated with them were then analyzed for similarity and overlap before equivalent codes were combined. During this process, the initial set of 138 codes was consolidated down to 129.

Phase 3: Searching for Themes

Rubin and Rubin (2011) explained that themes summarize and synthesize what participants express about a field, experience, or phenomenon, including details that provide insight into the nature of what is happening in the space, potential causes, and associated impacts. This goal of summarizing and synthesizing participants' statements while seeking details that explain the realities of teaching AI in CS classes guided my approach while seeking themes. Braun and Clarke (2006) describe this process of developing themes as poring over the codes and coded passages from phase two and experimenting with how various codes could connect together to create a comprehensive narrative or theme. During this code categorization process, when the data represented by multiple codes "look(ed) alike" and "(felt) alike" I collected them together into a category (Lincoln & Guba, 1985, p. 347). Even though research articles frequently suggest that themes simply emerged during a study, Saldaña (2021) points

out that themes do not merely jump out of data but instead are constructed by the researcher through their interpretation of the data. This was the case during the study's data analysis as the potential themes did not materialize on their own but only took shape through many iterations of attaching and reattaching codes together in seemingly meaningful ways. This resulted in the creation of the following eight categories or potential themes: 1) What AI and ML Are, 2) Why Students Should Learn AI, 3) AI Learning Goals, 4) Where AI Should Be Taught, 5) AI Pedagogy, 6) Unplugged AI Activities, 7) Challenges of Teaching AI, and 8) AI Curriculum. Five codes did not logically fit into any of these eight categories so were placed into a Miscellaneous category.

Phase 4: Reviewing Potential Themes

My analysis in the Reviewing Potential Themes phase focused on refining the previously identified potential themes. In this phase, Braun and Clarke (2006) suggest that candidate themes without sufficient data be pruned out to avoid analyzing less important themes that dilute the findings. They also suggest consolidating related themes in order to provide a more comprehensive whole as well as splitting themes that are too wide-ranging. Potential themes were first evaluated by reading through all of the interview excerpts associated with the themes. The data in each of the eight candidate themes was deemed to be cohesive, and collectively the eight themes seemed to tell the story of the entire data corpus. Next, the themes and their relationships to one another were considered together as a comprehensive thematic map (Braun & Clarke, 2012). This thematic map illustrated the connection between the AI Curriculum and Challenges of Teaching AI themes. Further analysis suggested that much of the data in

the AI Curriculum theme focused on negative aspects of curriculum. Additionally, of the eight potential themes, the fewest teachers (N=13) brought up AI Curriculum. With the emphasis that Focused Coding places on frequency of component codes and the overlap between the themes, I determined that the AI Curriculum theme should collapse into the Challenges of Teaching AI theme.

Phase 5: Defining and Naming Themes

After reviewing and refining the potential themes through the thematic map, the next phase was to finalize the themes in preparation for conducting the concluding analysis and report writing. To ensure that themes are not too wide-ranging, Clarke et al. (2015) encourage researchers to comb through all data associated with each theme and write a brief account that captures the substance of the theme. To do so, I summarized each theme's data with a narrative highlighting what the data illustrated and tied the theme's contents to the research questions. After completing the narratives, I reviewed their relationship with one another and identified a large amount of duplication between the Why Students Should Learn AI and AI Learning Goals themes. After realigning the subcategories inside the two themes, I refined the theme names so that they were descriptive, appropriate, and succinct enough to be readily referenced while writing the final report of the study's findings. This resulted in the following seven themes: 1) AI Understanding and Perspectives, 2) Importance of Students Learning AI, 3) AI Learning Goals, 4) AI within K-12 Curriculum, 5) AI Pedagogy, 6) Unplugged AI Activities, and 7) Challenges of Teaching AI.

Phase 6: Producing the Report

Braun and Clarke (2006) emphasize that the goal of the final phase of thematic analysis is to communicate the end result of the investigation or the messages that the data tell. While writing this narrative, I explained the details of the final themes and shared example passages from the teacher interviews to demonstrate the essence of what each theme represents. Effort was made to offer a comprehensive and detailed account of the context, research methods, and participants' perspectives in order to provide what Geertz (1973) describes as thick description.

Data Verification

When qualitative researchers both conduct interviews and then also perform the resulting analysis of the data collected as I did, there is increased risk of introducing bias to the study (Miles & Huberman, 1994). To help reduce this risk, Mason (2017) warned that researchers must be careful to emphasize participants' voices and take action to prevent their own thoughts from creeping into and contaminating the research findings. I employed member checking and intercoder reliability analysis in an effort to enhance the rigor, reliability, and validity of the study's data analysis.

Member Checking

Maxwell described the process of member checking as "systematically soliciting feedback about one's data and conclusions from the people you are studying" (1996, p. 94). To seek this feedback, I asked clarifying follow up questions during interviews whenever participants made statements that I thought could benefit from additional context. The goal was to verify my understanding of what teachers were communicating and provide them opportunities to expand and elaborate on their ideas. Subsequently, while reviewing interview passages during data analysis, ambiguous excerpts were

identified so clarification could be sought from the participants. I sent each teacher any interview passages that I was uncertain about to seek their interpretation. I contacted all 23 participants through the email addresses they provided and invited them to engage in this respondent validation. Ten teachers responded with specific insights on excerpts from their interviews, both clearing up misinterpretations and confirming conclusions drawn from the data analysis.

Intercoder Reliability

In order to improve the validity of the codebook and reliability of the data analysis, Lincoln and Guba (1985) advised qualitative researchers to have multiple researchers work independently to code the same passages and then compare the results. They contended that this process works to improve the consistent application of the codes used in data analysis. To do so, I identified a representative subset of quotes from the full dataset and trained another researcher on the application of the codebook. After coding the passages independently we compared results and found 78% percent intercoder agreement, approaching the 80% that Miles and Huberman (1994) recommend to demonstrate strong consistency between multiple coders' work. The Cohen's Kappa (Cohen, 1960) of the coding passages was 0.6471, suggesting substantial agreement based on Landis and Koch's (1977) benchmark of 0.6. Each researcher explained their decision-making and the codebook was revised to incorporate the new insights.

CHAPTER 4: RESULTS

Initial analysis of the teacher interview data led to 138 original codes. The codebook was reviewed for overlap and similarity and codes with comparable meanings or application were consolidated, resulting in 129 codes that were categorized into 22 subcategories representing the data. These subcategories were organized by similarity into seven themes that related to constructs being examined in this study. Specifically, the seven overarching themes were: 1) AI Understanding and Perspectives, 2) Importance of Students Learning AI, 3) AI Learning Goals, 4) AI within K-12 Curriculum, 5) AI Pedagogy, 6) Unplugged AI Activities, and 7) Challenges of Teaching AI.

AI and ML Understanding and Perspectives

One of the main themes that emerged from the interviews was related to teachers' understanding of what AI and ML involves. Teachers shared their conceptions of AI and ML in a variety of ways during the interviews that were categorized under three main subcategories: 1) perceptions of AI and ML and their abilities, 2) how AI and ML work, and 3) AI Types.

Perceptions of AI and ML and their Abilities

Seventeen teachers highlighted the mysterious and often opaque nature of AI. For instance, Mia suggested, "When I think of artificial intelligence, I think of like, very sci-fi stuff." Liam focused on how obscure and inexplicable the technology can seem, stating, "It [AI] looks like a magic trick." Teachers also commented on their students' concerns about how dangerous and threatening AI appear due to their often sensational portrayal in pop culture. Axel suggested the importance of helping students understand that, "At our level, it's just like, you know, it's [AI] not a movie. So it's not going to blow

up people and shoot them or kill them.” Another teacher commented on the importance of students recognizing that media depictions of AI are not always based in reality.

Damari explained, “I think a lot of people have like, an image that AI is like the Terminator or something. And it's a very different thing from like, what real life actually is.” In this way, participants pointed to their own and their students' conceptions of AI being informed and influenced by media portrayals of the technology, especially through sensational coverage. For example, Axel explained:

Some students think of AI as something that can take control of all computer systems and do things like launch nuclear missiles because they've seen this in movies like Terminator or Wargames or take over robots (from the movie I, Robot) and cause mayhem. This comes from what they see on social media, television, and movies. I think like, like most misconceptions, it is due to lack of knowledge on the subject.

Many of the teachers' (N=17) conceptions of AI's abilities were tied to the technology's capacity to replicate human intelligence. This was highlighted by one teacher, Damari, who put it as, “I think of it [AI] as trying to get computers to behave the way humans do.” Similarly, Penelope shared her understanding of AI as something that could complete tasks similarly to humans, stating, “Algorithms and machine models that are, that could, that can think intelligently, like humans.” However, one teacher, Itzel, pushed back against this common notion that technologies must be able to behave or perform like people do in order to be considered true AI. She felt that as the software becomes increasingly proficient, tying our conceptions of AI to humans' abilities could prove restrictive. Itzel stated:

Most people go with what you said about trying to replicate human intelligence. And, and really, I think, honestly, that's kind of limiting, right? What are some of the things that artificial systems could do that humans can't do? That's going beyond replicating our intelligence ... Why does artificial intelligence need to be limited to what human intelligence is capable of?

Additionally, many teachers (N=12) emphasized that the abilities of AI are driven by identifying human wants and serving to simplify their lives. For example, Queen said, "AI is a way to have a computer automate tasks"; while another teacher, Bella, shared her understanding that AI is a, "Piece of machinery that understands what you want, just by you giving it some information." Another teacher, Fatima, expanded on this notion and suggested that the software better meets users' needs the more they interact with it, allowing the software to learn from the users' actions. She shared:

Where it might respond to, you know, the input, but then based upon that input, if the customer is not satisfied, they might ask the same question in a different way. So then they're formulating a different answer. And the computer's essentially figuring out, well, maybe I should have done it this way to begin with.

While many expressed and demonstrated confidence in their understanding of AI and ML, five teachers had reservations about whether their current depth of understanding of the technologies would allow them to describe AI and ML to others. Tamika shared, "I don't know, I don't have that capacity to actually explain in a way that someone, like just a regular old friend of mine could understand." Jose emphasized that his conceptions are not fully formed yet, "Yeah, I don't know. I don't have like a set definition I go to anything yet." However, even teachers that doubted their ability to

articulate what AI and ML are during the interview suggested they are working to develop their understanding. For example, Winonah suggested that she is actively contemplating the technologies and developing her understanding, “I think, in my mind, and you know, I'm still kind of working through some of these ideas.” Another teacher, Ruby, expanded on this notion and observed that many educators will feel leery about their AI instruction until they have a deeper personal understanding. She remarked, “Teachers aren't going to want to teach about it if they're uncomfortable with it themselves. Like right now, I don't feel prepared to have anything other than a responsible use conversation without any further training or support.”

On the whole, teachers focused on the mysterious and nebulous complexion of AI as well as the common view of AI as a danger to humanity due to its depiction in science fiction movies and other media. Teachers largely deemed software's ability to mimic humans' capabilities as a defining characteristic of AI. This capacity not only defined AI in participants' eyes but also allowed the technology to benefit people as it could anticipate humans' desires and complete work for people. The bulk of the teachers interviewed were positive and optimistic about their understanding of AI and ML even though their knowledge was limited. Participants were generally confident in their grasp of AI and ML concepts. However, a handful of teachers doubted their comprehension of the technologies, with one highlighting that limited understanding of AI and ML might prevent many educators leading lessons on the full spectrum of AI topics.

How AI and ML Work

When discussing their understanding of how AI and ML work, teachers (N=16) frequently emphasized the technologies' use of algorithms that gives them sophisticated abilities. One of the teachers, Liam, explained, "It's more like an algorithm or a different classification of algorithms that can solve problems by simply just doing a ton of data analysis really, more than anything else." Some teachers expanded on this notion to include the importance of pattern-finding techniques that make use of mathematical principles. For example, Uri stated, "Machine learning is the, is a data analysis process that uses mathematical operations to identify patterns." Some of these depictions demonstrated more basic understanding of the technologies. For example, Cora summarized how ML works by explaining, "I mean, it's really, it's fancy math, that can take input and create output."

Ten teachers explained their conceptions of how AI and ML software is developed and honed. A frequent remark from teachers was about the importance of preparing and refining AI and ML through training datasets. Ruby explained, "It's the process of giving the computer training data so that it can then make choices on its own based on that information." Another teacher, Gemma, shared, "It [ML] has like a large database, and it starts trying to predict your outcome based on the patterns that it has seen in this database." Henry underscored the powerful and useful abilities that ML develops during this training process. He explained, "And it's sort of like datasets and training information that can give you meaningful responses to just about anything that you can throw at it." Some teachers also noted that after the initial training, AI and ML

systems can improve their performance even without further instructions from human developers; in essence, describing unsupervised learning. As Keshawn put it:

AI is something that can be iterative, and build upon itself with no additional human input, that's a big part of it, at least from my understanding, and that it takes outside input and builds within itself new generated responses and new generated outputs that are self serving.

However, at times teachers' conceptions of AI's and ML's functionality appeared to be driven by general ideas rather than a deep understanding of the technologies' underlying mechanisms. For example, Sebastian overestimated ML's capacity to grow and improve solely through user interactions, "So like, when we ask it [an ML system] a question, it probably is a question that it didn't have an answer to, but it found it out, because we asked a question." While supervised learning ML algorithms are able to improve their accuracy through training datasets of inputs and associated outputs cultivated and labeled by humans, the systems do not generally develop new capabilities merely through interacting with users (Goodfellow et al., 2016). Another teacher explained that the seminal ML program is written by a human but mistakenly asserted that the underlying logic is rewritten by the program itself. Fatima described the process as, "The machine, the computer is, could possibly [be] rewriting its own algorithm too but an initial algorithm was put together to compare research data, and then put together an answer." This conception is likely inspired by the reinforcement learning ML approach where the agent, or algorithm, is designed to learn through independent trial-and-error, continuously being rewarded for actions that move toward a predetermined end goal such as safely navigating an autonomous car through traffic.

The previous lessons learned inform the agent's current decision making as it continues to hone its choices and ingest new information. Fatima's conception of this type of ML is largely correct; however, the core reinforcement learning algorithm is not rewritten by the system. Instead, the agent continuously updates its decision making policy based on what it learns from its trials and errors (Sutton & Barto, 2018).

Altogether, teachers primarily conceived of AI and ML as algorithms or advanced mathematical procedures which identify patterns in sample data and extend those patterns to create the desired output for the user. Teachers' conceptions of the AI training process concentrated on providing human-written models with vast amounts of training data in order to educate the system to provide useful outputs to users. The quality of that output continues to improve the more users interact with it since these interactions create more data points for the system to analyze. However, some participants overestimated AI's current ability to reengineer itself.

Types of AI

Only 11 teachers provided example types of AI, all of which were forms of ML including classification models (N=6), regression models (N=5), neural networks (N=4), natural language processing (N=4), and principal component analysis (N=1). Teachers most frequently mentioned image classifiers as a form of AI they have encountered. Damari described image classifiers as, "So the obvious answer is like, let's say you have images and you're trying to classify them as one type or another type." Winonah explained, "So machine learning is sort of the system that enables image recognition." Teachers also commonly referenced regression models as a common type of AI. Jose, when describing ML models explained that developers, "Train a model to basically

either make decisions or whatever else you want that give you some type of analysis, right, like regression analysis or anything like that.” When discussing types of AI, several teachers also pointed to natural language processing. As Penelope explained it, “For example, something like natural language processing, which would be something like a chatbot where you could chat back and forth with a program online that could answer your questions.” Neural networks were also described as a form of AI, though the teachers that talked about them generally did not feel confident about their details. For example, Itzel shared the following while explaining a tool she uses with students to illustrate neural networks:

I say, “Here's the Google Teachable Machine, train it. This is what a neural network looks like. And it's these orange and blue dots. And nobody knows what the heck is going on. And we're adding neurons and we don't know what any of this means.” And I tell them, that's okay.

Other teachers occasionally alluded to the fact that there are further forms of AI in addition to ML but did not share examples of specific types. In the words of Gemma, “And that's, you know, how the machine learning process works is that it tries to find patterns. And that's just one form of AI.” Mia acknowledged that she applies the names AI and ML synonymously to the technologies. She explained, “I feel like I use the terms [AI and ML] interchangeably. I don't know if that's accurate, though. I guess, because my understanding of artificial intelligence is that I guess currently, I think it is limited by machine learning.”

In summary, the majority of teachers did not identify types of AI but those who did most commonly shared classification models, regression models, neural networks,

natural language processing, and principal component analysis. While several teachers articulated that ML is one of multiple forms of AI, others felt that AI and ML are one in the same and that the technologies are identical.

Importance of Students Learning AI

Another theme established during data analysis was why it is important for 9-12th grade CS students to learn AI. The arguments for teaching AI to K-12 students fell into three subcategories: 1) AI knowledge is necessary for today and future, 2) citizenship, and 3) awareness of bias and potential harms.

AI Knowledge is Necessary for Today and Future

One of the main arguments teachers used for developing competence with AI was that it is important not only for students' lives right now but will also be increasingly critical for their future (N=20). Elliott put it plainly, stating, "The reality is, is that it's [AI] here and it's here to stay." Because AI has become so ubiquitous and is becoming even more ingrained in our lives, these participants felt that their students' future success depends on a knowledge of the systems. Keshawn shared, "I feel like it's absolutely important to teach the upcoming generation of kids who are probably going to be around our [teachers'] age when AI is ubiquitous and everyone's using it."

Participants also emphasized the importance of AI knowledge for students' career prospects. For example, Cora stated "There'll be a lot of career opportunities, I think, too, if you have some level of literacy there." If AI knowledge will help students excel in the future, teachers felt that schools owe it to learners to provide opportunities to develop AI knowledge and skill. Liam explained it this way:

The job market that these kids are going to be in is going to, is going to be one that leverages AI and AI tools for, you know, various different things. So that you can't see we're doing our students a disservice if we're putting them out there into the world of you know, being an adult and trying to find a good job for yourself and we haven't taught them anything about it.

As high school CS teachers, some participants (N=3) observed that learning about the technical mechanics of AI provides a valuable opportunity for CS students to apply their coding skills. Students in high school level CS courses generally have budding programming skills and the chance to write software that incorporates AI can be motivating and also serve to strengthen their skills. One teacher, Ruby, told of her CS students' enthusiasm for writing software that incorporates AI. She explained, "I have kids that are already interested in like, how that works on the development side, like how they could create their own AI tool." Another teacher, Violet, highlighted how learning to write programs that employ AI would allow students to use their programming know-how and reinforce the skills they are learning in class. She pointed out:

The other part to it that I think would be so beneficial is that it is an application of the skills that they're learning in the coding class. So they're doing these coding classes, they're learning coding stuff, and they're working on stuff and you kind of see the benefits, but take that into an AI thing, it's like, 'Okay, now you can take those skills, and actually develop something that's being used.' I think that would be a great thing for applying the skills that they have been developing.

On the whole, participants felt that learning about AI is necessary for students' present and future. They argued that knowledge about tools will help students prepare for career success, and utilize their nascent programming abilities in CS courses.

Citizenship

A significant majority of the teachers (N=21) felt that it is crucial for students to develop appropriate citizenship in the contemporary AI era. These teachers expressed that students would need a level of AI literacy in order to be prepared to understand and contribute to our society. As Cora explained, "They're gonna have to be somewhat AI literate, when they go out into the world, even if they are just a regular citizen, I think that would be valuable." Other teachers emphasized that while not every student needs to deeply understand all of the technical details of the technology, there is value in learners comprehending how it works at a basic level. For instance, Uri pointed out:

At the level that we understand from working, you know, the level that we understand, you know, the way electricity powers our homes, and the way we, you know, understand that there are laws in the country, to me AI is sort of going to be a foundational aspect, like the way we understood, like, we all could probably generally explain how cars are put together on an assembly line, even a lot of the kids could, and no reason they shouldn't be able to do that with AI.

Many teachers (N=19) also pointed to citizenship skills allowing students to develop the ability to differentiate between proper and improper applications of AI. Participants emphasized the importance of providing opportunities for students to consider the responsible use of AI. When sharing about the significance of these learning opportunities for students, Ruby suggested, "They need to know what's good

for and acceptable and responsible use of it [AI]. So I think we would honestly be doing kids a disservice if we didn't have meaningful conversations about that.” Ruby went on to discuss that the alternative to formal instruction about acceptable AI use is for students to make their own way with no direction from trusted adults. As she put it, “I wouldn't feel very comfortable with them just finding it [AI tools] on their own and not having any guidance at all. We have enough problems with kids finding things on the internet and having no guidance.”

Teachers emphasized that no one is immune from the impacts of AI and helping students recognize the technology's effects on their lives and lives of others can help them develop a fuller digital citizenship. Bella explained, “I think it's important, I think it's, it's [AI] a technology that's out there, it's being used, we're all victims of it, whether we know it or not. So to learn it would only make us better digital citizens.” A specific impact of AI that teachers mentioned is the lack of privacy due to the technology's collection of user data. For instance, Penelope underscored the importance of students recognizing how AI harvests data while users have little control over how that information is used or shared. She stressed, “I would want students to understand just how much data is collected with every single thing that they do, how that data can be used responsibly or misused.” Keshawn also felt that given the lack of privacy, it is critical for students to evaluate how they interact with the technology and the sensitivity of the data they reveal to the software. He asked:

What information would be responsible to give to an AI with what do you know what they do with that information? Some of that more like, you know, your internet footprints forever kind of concept, but just in a different way? What does

it mean to have your information put in a black box that probably no human touches, but it's still utilized by that proprietor of the AI?

Further, teachers clarified that learning about AI will not just allow students to examine the technology's damaging impacts but also help students see their role in reversing the biases and improving the software's design. For example, Queen shared her feelings that students should receive instruction on AI ethics to help them recognize the opportunity they have to limit the damage that bias in the technology causes. She stated:

Ethics, definitely understanding how AI is currently biased. And that it will be, it will be their responsibility to make sure that it doesn't get worse. And for that, they have to understand why it is biased, how it's biased.

In sum, teachers shared about the importance of students acquiring AI citizenship skills in order to allow them to be productive citizens who are prepared to engage in discourse about the merits of the technology. Participants also felt these citizenship skills will help students make informed decisions about how AI should be applied responsibly in their lives and throughout society including identifying and correcting the use of AI when misapplied in damaging ways.

Awareness of Bias and Potential Harms in Technology Design

When explaining why it is important for K-12 students to learn about AI, a number of teachers (N=18) spoke about how an understanding of AI helps students' awareness of the technology's potential harms. They argued that dangerous biases are inherent in AI systems but are typically working beneath the surface. As one teacher, Mia, put it:

And there's, there's nothing that is without bias. And so we can't blindly trust an algorithm, even though I think a lot of people think you can. But we just, it's unreasonable to trust an algorithm. Because the person, the people that are involved have biases, and usually their unseen biases are unacknowledged.

Teachers stressed that because these biases are not typically apparent to users of the technology, it is vital to provide students well-designed learning activities to help raise the visibility of the potentially negative impacts of AI. One participant, Violet, emphasized the importance of providing students with opportunities to reflect on the sources of bias and the influence that a diverse range of perspectives can have on the AI design process and level of bias in the final product. She explained:

I think bias is really important. I think being to being able to critically read about things, and to, and if you were involved, to be able to critically think about okay, 'What are possible biases that are creeping in and where are we not expanding the pool of information that's being used and the pool of people that are being used?'

Teachers also wanted students to recognize that because the programs are written by people, they reflect those humans' biases. When teaching about AI, Gemma discussed how she reminds her students that any system "Can only be as good as the person who creates it." Teachers pointed out that it is crucial for students to comprehend that bias can slip into AI even when the designers and engineers have good intentions if they do not scrutinize their methods for potential lapses. For example, Henry shared, "There can be potential problems with machine learning, depending on like, you know, it's always only as good as its training set." Similarly, Winonah explained

that she works to help students see that “The quality that you see as a user of AI is based off of what's going in.” Keyshawn also underscored the importance of hands-on learning activities that help students experience the potential for biased or tainted training data to impact the resulting AI that students create.

That's what I think would be the biggest thing that they'll [students] pull out of that [AI class projects], that GIGO right, garbage in garbage out. That's only, you're [the quality of AI is] only as good as the information you put in.”

Further, Gemma shared that students should recognize that the quality and danger of AI mirror the quality of the design process and training data curated by humans. She explained:

I hope they leave with the understanding that junk in junk out basically, like, you know, it's, it's how well these things are going to perform is totally totally up to how well they're designed and what, what data is being used and the engineer or our computer scientists that's developing it.

Overall, teachers stressed the importance of students learning about AI to better understand the potential harms that the technology could inflict. They argued that this could help learners identify how humans unintentionally introduce biases into AI through the design and training processes. Participants emphasized that better understanding AI will help students see the crucial role of diverse, representative design teams and critical examination of potential training data.

AI Learning Goals

Teachers also detailed their views on the primary learning objectives for K-12 AI learning experiences. The 9-12 grade CS teachers' responses were organized into the following three subcategories: 1) AI literacy, 2) AI ethics, and 3) AI for learning.

AI Literacy

Almost all of the teachers interviewed (N=22) advocated for K-12 instruction to include objectives to demystify AI and increase students' AI literacy. They discussed how mysterious AI can seem to users, including students, without an understanding of how they work. For example, Liam explained, "It [AI] looks just like, like a black box, the black mystery box that does some strange thing, unless you actually know what's really going on." Participants felt that as society becomes more reliant on the technology, it is important that students have opportunities to deepen their awareness of how it works. As Elliott put it, in these lessons students will be "Getting a better understanding of the devices that they use all the time" and "I want to demystify the whole thing for them." Mia explained that as she considers how to demystify AI she considers, "How do I make it so that it's [AI] not like, you know, a black box."

Further, teachers felt that stripping away the mystery of AI can allow students to better consider the technology's most productive applications. Keshawn pointed out, "People need to understand how this [AI] works and what it does and what it can do for us and how it can be used." This understanding is a building block of AI literacy and can unlock an ability to question and evaluate the technology once students know how they work at even a basic level. Damari shared:

I think AI literacy is a very good thing to add into that group of ideas. And a lot of ways like, I guess I think of it a little bit like when you're reading the news, how do you figure out whether it's a good source or a bad source sort of thing? Right? You have an AI literacy of like, "Okay, what did these things actually involve? How are they learning? What are the impacts of that on what I'm doing?" All of a sudden, Amazon keeps showing me these ads for something I just looked up. "How is that actually working?"

Nine participants expressed that teachers can empower students by strengthening their AI literacy. If AI instruction can remove the barrier of mystery from the systems, students are enabled to affect change. One teacher, Liam, claimed that developing an understanding and comfort with how the technology functions allows students to thoughtfully apply it in the appropriate situations. He shared, "Knowing how the solutions work and how the problems work is knowing how to pick the right solution for the right problem." Other participants felt that this empowerment from AI instruction could extend past students applying the technology and would inspire them to recognize the impact that they can have on the systems themselves. Itzel explained:

So when they start understanding, this [AI] is not magic, and there's something that can be done, or they have a level of influence and can talk about it. That helps empower them to understand the situation and maybe what they can do about it.

Queen emphasized how critical enabling our youth could become to our future if the technology becomes further tainted at the same time that it becomes more deeply

ingrained in society. She asserted that tomorrow's solutions rely on providing appropriate AI instruction to the future generation today. She said:

Even if our generation ruins the whole thing, you know, they, they still have it in their hands to, to fix some stuff, and to start now, to try to make a difference. And if they don't understand how it works, if it's, if it's magic, then they're not going to be able to do anything about it, it's going to continue being dangerous.

Teachers also spoke strongly of the importance of students using AI tools in order to provide them with exposure to the technology and help demystify them. This guided use can also lead to many constructive discussions about the technology. As Itzel explained, "If they're interacting with the tools, then they can learn about them. And they're not necessarily going to be learning about them at a nitty gritty level. But 'What is it? When has it been frustrating?'" Another participant, Nakeisha, felt that first-hand knowledge of tools is a critical component of AI literacy. She shared, "By the time a student reaches high school, they should have a good definition of what AI is. And they should know different tools. That's pretty much the baseline of literacy of AI."

In essence, participants lobbied for students to receive AI instruction because demystifying the technology could unlock students' ability to better utilize it. Participants shared the value of AI literacy to unlock the nebulous black box to understand how AI functions. Teachers also argued that in addition to better application of AI tools, a deeper understanding of how the technology works could embolden students to fight the dangers it poses. Adequate AI literacy can help students appreciate both the abilities and limitations of the technology.

AI Ethics

Given the potential for harm from AI, the majority of teachers (N=20) felt that ethics should be a key component of K-12 instruction. For example, Elliott shared, “It would be advantageous to get instruction to students about how to use it [AI] ethically.” Teachers largely felt that students should learn about the ways humans influence AI and how AI influences humans. Even though the inner workings of the software can seem opaque to students, these participants wanted to ensure that learners recognize that AI is created and shaped by human hands. Violet, for instance, spelled out that “The whole point of artificial intelligence is writing programs, which so often is kind of not seen by the general public, but it's programs in the computer, that are written to have the computer generate its own information.” Queen also emphasized human involvement in the development in AI, in particular, training the systems. She said, “I think that's the important part that they should understand that, that the training of the AI is still done by humans.”

As today's students will be tomorrow's AI systems developers, teachers also stressed the critical nature of discussing responsible design during these formative years. Teachers emphasized that without it, AI is prone to a wide array of dangerous, unintended consequences. Itzel explained it this way:

We talk about some of the consequences, so not only are there huge benefits, but the consequences of a lot of times are unintended. People start out with, with wonderful intentions, ‘Oh, this is going to help people so much, this is going to be so wonderful.’ And then really, really bad things happen. ‘Why does that happen?’ And we talk about a lot of the issues in tech is a big rush to the new

and shiny and to do the cool thing. And slap on some ethics, maybe some security, maybe some privacy after the fact and it never works very well.

In order to combat these unintended consequences, teachers encouraged an emphasis on the importance of representation in the design and training phases. One teacher, Damari, provided the example of when developers might reflect and say, “‘We accidentally didn’t think about this other group that we’re biasing against, without really meaning to.’ Like you put together this well-meaning algorithm, but you didn’t look at your data to see if it had bias in it.” As a means to minimize the introduction of bias, Mia emphasized that those designing the technology must have a focus on the diverse user base that will eventually use the product. She asserted:

The most important part of any development is the people who are going to be using the development and making sure that their concerns, their, not their concerns, but making sure their humanity is like being factored into the development.

In addition to the potential harm that comes from human influence on AI, teachers also emphasized that students should recognize their potential to influence the technology for good. Keshawn encouraged instruction to help students consider, “What ways can you impact or change the AIs that are out there?” Similarly, Nakeisha shared optimism that learners will be able to bring solutions to harmful AI by coming to it with a fresh perspective. She stated, “My hope is always that students will maybe fix something that is broken.” Itzel suggested that students could be empowered not only to influence but also to build new AI when they understand the technology. She said:

One of the things I think is super important is that the kids understand that they're, these machines are not magic, that their phones are not magic, that they could build anything in their phone. Should they want to, they can make that happen, if that's something that they're interested in, and that they know enough about it to know, to be able to influence it and potentially create with it.

On the whole, teachers felt that K-12 AI instruction should include ethics topics to help students reflect on how humans shape AI. This can help them better see how the quality of AI systems is dependent on the quality of the design and training performed by engineers. Teachers emphasized responsible design of AI and the value of representation in the design process. Participants also suggested that students have opportunities to contemplate how AI influences humans, particularly how biased systems harm individuals and how students can improve AI software for the better.

AI for Learning

Another 11 teachers advocated for K-12 schools to teach students to see AI's potential for learning in a variety of areas. Teachers praised the support AI can provide students as they conduct research. Elliott said, "This thing [ChatGPT] is actually really good at finding sources ... you can definitely use it to find sources and then go do your own research." Similarly, Keshawn praised, "This [large language models] is something that's helpful for your [students'] research." In addition, teachers spoke of AI's capacity to support students and help them develop wide-ranging skills. For example, Winonah told of a Language Arts teacher who encourages students to use AI as a real-time writing tutor. She said:

He taught his students how to use ChatGPT to improve their writing. So he taught them how to write a basic essay, you know, whatever they're doing, and then go into ChatGPT and say, ask a query like 'What is the what is the tone?' I don't know 'tone' is the right word. 'What is the feeling of this essay?' Maybe or 'Can you score the grammar on this essay?' or that sort of thing, and then he taught them how to look through what ChatGPT has, and then improve the writing so they can use it for immediate feedback.

Participants also detailed the benefits of AI tools in the CS classroom. Penelope explained that she encourages her CS students to use AI tools as they program to help them see how they can save time and effort. She said:

There are AI tools now being built into the IDE to help with, to help programmers to develop algorithms, you know, that maybe they don't have to write from scratch, instead of using a library or an API. Now, they [students] could use an AI tool built within the IDE.

Liam also spoke about the potential of using AI in CS courses to help students understand new programming structures and approaches. He explained:

[Chat]GPT will give you like, step by step, play by play, puts comments and everything for you, and kind of explains how it works. So as, so I see it, as I do tell them that I see it as a learning tool.

While teachers generally discussed the value of learning through AI tools, participants also stressed that students should have opportunities to learn that AI may not be the most appropriate and trustworthy tool for every task. Sebastian explained how students can learn with AI, "They [AI] can solve math problems, they can solve

science problems, they can cure cancer, they can find a solution to a mathematical theory.” However, he went on to express “I want students to understand the things that it [AI] can't solve, I want students to know that an AI that learns and gives information and all that has its own limitations. It's not the answer to everything.” Teachers stressed that learning opportunities for students should aim to help students see that AI systems are not all-knowing, universal solutions that are able to solve all problems without human intervention. Liam explained his instructional approach as, “We sort of like to study kind of the limitations of it [AI], and how it's not a, it's not a cure all.” Instead of competing with or becoming overly reliant on AI, teachers pushed for students to receive instruction that could help them come to view the technologies as partners that help us solve problems more effectively. Jose explained:

The people that are successful in implementing AI technologies in their businesses are the ones that are taking the approach of working with AI, as opposed to trying to use AI to replace tasks, right? So we use it to give us knowledge to help us make decisions and stuff like that.

In order for students to meaningfully harness AI for learning, teachers argued that they need to develop the skill to evaluate the plausibility of AI-generated products. Keshawn explained the type of thought process he would like to see his students exhibit, “I would even encourage them to say, ‘Cool, here are some AI generated solutions for this problem. Which one of these makes the most sense? If one of them does, does it even fit our problem?’” In the end, a well-rounded understanding of AI will allow students to judge which potential tools might help solve a variety of problems and identify the best fit. As Keshawn explained, “I think the same thing could be said of AI as

when to know when to use the tool or when to use your head and when to put the two together.”

Complicating instruction about AI for learning is the fact that students can use the tools in a way that many teachers deem cheating or plagiarism that in effect prevent students from learning. Penelope gave voice to this when she shared, “There's this big looming elephant in the room of, ‘Okay, these tools are being used to cheat.’” When considering how they can proactively help students discover the academic acceptable use boundaries, teachers largely shared Fatima’s sentiment that schools should, “Teach them how to use it [AI tools] appropriately, rather than cheating with it.” A common thread among participants was that in order for teachers to help students recognize how to learn through the technology, schools must allow students to use AI rather than universally banning the tools for fear of inappropriate use. Liam encouraged teachers to help students build their own agency about learning rather than hard and fast rules about the use of AI tools on all assignments. He shared:

You could kind of gently give them the right idea about how you shouldn't really be using this [AI tools] for all of your assignments. Then, because then you're not really learning anything, you know, these are the conversations we'd have to have.

Much of teachers’ concern about cheating centered on generative AI’s ability to instantly create writing and other works that closely resemble what students create themselves. Axel explained, “The computer generates all their stuff for them and, and other teachers are worried that you know, kids are gonna copy and just use it and turn it in as their own.” Penelope explained her approach as, “With my students, I have asked

them to check into it [AI], you know, because I want them to be aware and to understand I don't want to put blinders on their eyes. I want them to use it responsibly.”

Liam also encouraged teachers to reframe these cheating concerns as learning opportunities to engage students about responsible use. He advocated:

You could talk to them about the limitations, you could kind of gently give them the right idea about how you shouldn't really be using this for all of your assignment. Then because then you're not really learning anything, you know, these are the conversations we'd have to have.

Discussion about AI for learning centered around students' capacity to use the technology as a means for discovery. Teachers also wanted students to see that AI's abilities are not always a clean fit to directly replace human intelligence and decision making. Many teachers suggested that students should come to appreciate that humans can harness the technology to enhance their abilities after careful consideration of the software's strengths and weaknesses. This deliberation should also take responsible use into consideration and evaluate whether the application of AI is merely plagiarism or informing students' own work.

AI within K-12 Curriculum

Another theme that emerged from the data was where and when K-12 CS teachers felt AI should be taught in K-12 schools. Teachers' opinions varied but generally fell within the following two main subcategories: 1) AI literacy for all and 2) AI in elective high school CS courses. It is important to note that participants approached this conversation with a wide range of experience teaching AI in their high school CS courses. Six teachers had taught a standalone AI or ML elective CS course, 12 had

incorporated AI topics into their existing CS courses, and five had not yet taught AI concepts in their courses.

AI Literacy for All

Nineteen teachers spoke about the merits of all students receiving instruction on AI. These teachers largely felt this approach would be advantageous so that all students would have the opportunity to develop their AI literacy including those who would not opt in to learn about the field through an elective course. For example, Elliott argued:

I think since this [AI] is so central to the technology space that we're in now, it would be good if they [all students] got just a little bit of experience, if everyone got a little bit of experience in it.

Another teacher, Cora, felt strongly about the value in introducing all students to AI. She shared, "I think at least exposure is going to be super important, you know, maybe a month or something of exposure." Similarly, Sebastian felt that not every student needs to learn all of the technical details of AI, but requiring the topic for all learners would raise their AI literacy. He explained:

Oh, I think it's super necessary [for all students to be taught about AI], if at least just to expose them to it. I think it's something I'd say it's a topic that needs to be brought up now because it's [AI] becoming very popular.

Elliott further explained this line of thinking that every student deserves an exposure to the field to help them build literacy, and then they can decide for themselves whether they are interested in learning more. He said:

I think yeah, having, having the information for, for all students would be the best approach. Yeah, not every student is going to be technically like inclined to go

down the technical path. Although I think, I think there's advantages to, to it, just integrating pieces of it for, for all students.

Some teachers felt that students would benefit from learning about AI in the context of other required subjects rather than a required stand-alone AI or CS course. For example, Bella suggested, "It would be fun to learn how to integrate it [AI] into other courses." Another teacher, Penelope, also endorsed this approach. She said, "I definitely think that it would be an ideal world if it [AI] was addressed in other subjects." Orin also felt that embedding AI instruction in other subjects, particularly mathematics, would benefit students. He shared, "It would be a good exposure and if I, you know, I guess if it could be linked into the curriculum, especially in math."

Teachers pointed out that one way to ensure that all students learn AI is to introduce content standards related to the technology as standards largely determine what is taught in K-12 schools. Axel stated, "It's probably needed, does need to be talked about becoming more into the standards than it probably already is." Keshawn called attention to the fact that the standards-setting process can be methodical and slow to respond to rapidly-evolving technologies such as AI. He shared, "As of right now, it's [AI] non-existent, at least in our standards, and <State> is not known to be the fastest adopting technological standards."

Many teachers suggested that schools offer all students AI literacy instruction before they enter high school. These participants conveyed that K-8 lessons could offer foundational knowledge to students at an early age. For example, Elliott felt that, "Fifth, sixth grade could be a good, a good place to start. And then from there up." Ruby felt that waiting until high school for AI instruction skips valuable opportunities to equip

adolescents with digital literacy skills. She expressed, “Knowing especially how 10, 11, 12 year olds get into other sorts of trouble on the internet, and probably needs to start middle school level, maybe just with ‘Here, here's what it is.’” Likewise, Axel argued that, “It'd be good in the lower levels, to just have an introductory ‘This is what machine learning is, this is what AI is, this is generally what's happening when we do that.’” Axel further discussed that all elementary students should be exposed to CS which would help inform their elective choices related to AI in secondary school. He said:

So I think in the lower levels, if there was more awareness of what programming is, computer science, that kind of thing, so that as they get older, they can make better decisions in middle school to, for their electives that they take there.

Several teachers cautioned that K-8 students' background knowledge and abilities might be insufficient to learn how to develop AI tools independently. They felt that instead, instruction focused on AI literacy may be more suitable. One teacher, Damari, explained the importance of presenting AI learning activities that are developmentally appropriate for younger learners and push their understanding forward. He shared:

There is a possibility that junior high students could use partially built or sandbox environments for creating and using AI since they don't have the statistics, the calculus, the necessary math skills to create AI systems themselves. That could, it would be a reasonable place to start exploring those systems, though. The trick is to make sure they are actually learning something useful and not just playing with a chatbot or trying to find ways to get the AI to write their assignments for them.

Multiple participants pushed back on the possibility of teaching K-8 students to write AI due to their lack of programming experience and readiness for sophisticated coding concepts. For example, Axel shared the frustration students can experience when they may not be developmentally ready to code. He shared:

I just think the development of the students, as they get older, they just become a little bit more aware. Of the kids in the lower levels, I'm not sure they're ready to handle programming. I mean, we have a lot of the eighth graders that I taught last year, there are a few in there that wanted to learn how to program but then they were just a little bit lost on the whole how to program and why it works and it's fairly frustrating for a lot of them. So I think that that end is why I think it's better for later on. Because AI is pretty complex.

Gemma suggested avoiding instruction for younger students that requires programming experience. While this may limit AI learning activities, lessons for inexperienced learners could focus on building AI literacy. She explained:

Well, as far as the younger kids, I don't think it'd be easy to have them do a lot of hands on because it uses, it uses so much, uses a lot of programming skills that they may not have yet. So I would want them to mainly have that literacy at a younger ages.

Some teachers also believed that K-8 students might not have the level of maturity and digital literacy necessary to engage in discussions about the implications of AI yet. Damari pointed out:

They should probably know the consequences of the things that are happening in their digital lives. From the creation standpoint, I think it probably comes later on,

it probably does, you probably don't get AI until you make sure you have a stronger skill set in order to understand what the creation of these things even means.

In essence, many teachers felt that every student should receive instruction about AI. They promoted this universal approach as a way to ensure that all students receive exposure to the field and build their AI literacy. Some teachers suggested teaching AI in required core courses, and multiple participants championed building AI concepts into content standards to ensure it is taught. A group of teachers advised that early exposure in grades K-8 could help students develop core understanding of AI from a young age while others had reservations about younger students' maturity levels and background knowledge.

Elective High School Course

When considering where AI should be taught in K-12 schools, a number of participants (N=17) discussed elective high school courses. Teachers argued that the high school years are a logical time for students to explore how AI functions. If that introduction sparks an interest for students, they can then choose to take an elective CS course to learn more about the technical aspects of the systems. Winonah shared:

Developmentally and interest wise high school would make sense to then jump into some of the more technical pieces [of AI], maybe for all students to have a basic understanding or for those who really are interested in, in the details that could come in a CS course or somewhere.

Several participants argued that because not all students are interested in AI they should not all be required to learn about the technology. These teachers suggested

that indifferent students would gain little from a required AI course. For example, Axel shared that students who do not elect to take his current CS courses but are forced to enroll do not always benefit from the instruction due to their lack of interest. He said:

I get students now who don't actually opt in, they're just put into some of my classes, and they have very little interest at all in any of that kind of stuff, so it doesn't meet their needs at all.

Other teachers raised similar concerns. One teacher, Orin, shared that his own children did not enjoy their experience in required technology courses. He described, “My kids were kind of forced into taking a coding course, they didn't like it.” Similarly, Sebastian emphasized that every student is different. While all students should have the opportunity to learn about the basics of AI, those that are not drawn to the field should not be required to learn more about the mechanics of the technology. He explained:

I'm thinking of student, of my students and some of them will not be interested in the technical side of it, like, some will, because some are intuitive, though, they'll want to know how it's built, how it's created, how it's programmed to learn from itself and from other things. But I think mainly as a, as a kind of like a general thing. I think it's mainly the exposure and the ethics to it. Because not every student is looking for that technical knowledge.

Teachers also discussed the challenge of incorporating AI into existing CS elective courses that make use of rigid curriculum. Courses, such as Advanced Placement CS A (AP CSA) or International Baccalaureate CS (IB CS), often have specific topics that need to be covered in order to prepare students for end of course assessments. Even though these courses would be a logical place to teach about AI,

the topics are not currently included in the course content. Cora explained how updating the AP course content to include AI would benefit students. She explained:

I mean, foundationally, I guess it would be nice if like AP, actually could be a little more nimble and actually see, you know, yes, it's important to teach all these things about Java, but like there's other important pieces that are being left out entirely. And maybe they ought to update a little.

In summary, a number of teachers recommended that AI be taught in high school elective classes rather than required courses. Many of them argued that this would allow students a choice whether they learn about the field rather than forcing all students to do so. Teachers also spoke about the challenges of incorporating AI into elective courses, especially CS classes that rarely have space for additional topics.

AI Pedagogy

Teachers frequently mentioned a variety of teaching techniques they felt are important when providing instruction on AI concepts. The teaching approaches fell into the following four subcategories: 1) student centered pedagogy, 2) teacher led pedagogy, 3) student tools for instruction, and 4) class communication norms.

Student Centered Pedagogy

Thirteen teachers shared student centered pedagogies they believe important to successful AI instruction. These participants described a range of teaching approaches that allow students to engage with concepts in order to actively make meaning. Damari explained that he chooses not to deliver instruction through lecture but instead provides opportunities for students to engage with AI course content individually. He explained:

So I don't believe in standing and lecturing at kids ever. I feel like that mode of teaching doesn't work very well, especially for something that's a very practical skill that each person kind of has an individual relationship with as they start to learn how to work with it.

Mia avoids this teacher-centered approach as well, simply stating, "I really don't like to lecture." Instead, many teachers spoke of the benefit of hands-on approaches to learning. For example, Violet encouraged AI learning opportunities to allow students to program in order to develop an understanding of how the technology really works. She shared:

Giving them an opportunity to do some coding along those ways [implementing AI algorithms], I think would be really great. I think that would really introduce them to the whole concept of what it is from the technical side of things.

Teachers frequently spoke of projects as a positive way to allow students to engage with AI. Uri explained his view that projects invite real-world uncertainty into the learning environment and also permit the support structure to aid students when necessary. He described:

My belief is that by working through a project, working through ambiguity, I use that word a lot with my students from the beginning, being curious and using ambiguity to learn. And so projects, I think, are the greatest way for them to kind of confront that ambiguity, often with somebody to support them up here, and to have enough time for each person to work through those projects.

Inviting this uncertainty, or ambiguity, into the classroom through projects can empower students to overcome challenges through trial and error. As Itzel put it:

“Sometimes just experimenting and playing with different things, and just being, having fun with the frustration has been helpful.” Nakeisha explained how a project-based approach to instruction enables her students to take control of their own learning. She said, “I love to give students project-based. I'm really huge into allowing them to be investigators, explorers of the information.” Another teacher, Jose, conveyed how projects benefit students by allowing them to put their skills to use in meaningful ways. He reported, “It's [projects are] like actually applying the knowledge to like a real life situation essentially.”

Teachers also detailed other pedagogies they find fitting for AI instruction because they provide autonomy to students. One teacher, Damari, shared that he uses a flipped classroom instructional approach where he pre-records video instruction for each lesson and allows students to work through lesson materials and projects independently. He stated that this allows each of his students to access lessons in a way that works best for them and simultaneously frees him to support each student and tailor instruction to their needs. He said:

The reason I do the flipped classroom and like, try to individualize the lessons closer to each individual person is because each person is different. And so I'm trying to find the issues they have as fast as I can and make sure that we find a way that they can learn as quickly as they can for their particular learning style, and their particular background.

Two teachers also detailed their affinity for Socratic-style, student-led discussion. Winonah shared her view of having students actively explore AI concepts through

open-ended conversations with their classmates, rather than traditional, teacher-led question-and-answer type discussions. She explained:

Socratic seminar type discussions, and having some of those kind of student driven, open-ended top conversations around AI and maybe it would be helpful to have like, some case studies or some like live examples that we could run real quick.

Uri also spoke about how Socratic method type discussions would allow students to analyze how AI is used in real world applications. He shared, “We could have sort of like a Socratic class about AI. And where we would, you know, maybe kind of build maybe, so examine some case studies about implementations of AI.”

When sharing about student centered pedagogies for teaching AI, teachers spoke about the efficacy of no-lecture teaching, hands-on programming, project-based learning, flipped classroom-style video instruction, and student-led Socratic method discussions.

Teacher Led Pedagogy

In addition to student centered teaching methods, many teachers (N=12) explained the merits of a variety of teacher led pedagogies for teaching AI concepts. One common teaching technique described was for the teacher to introduce current events and societal issues related to AI. Ruby shared her approach to draw attention to AI in her classroom “Using some of the current news cycle.” Similarly, Orin reported that he is quick to introduce current events involving AI into his courses. He said, “When news comes up I bring it to them as kind of an enrichment thing.” Mia, when speaking about how AI can lead to dangerous outcomes if not carefully designed, described the

importance of giving actual contemporary examples to help illustrate potential issues.

She detailed:

Giving students time and space and examples of how that happens, not just in technology, but kind of just in life. Because I think that can be very, very hard for students to conceptualize. I think examples are probably like the most useful tool.

Beyond open-ended, student driven Socratic seminar type discussions, teachers also described employing teacher directed discussions focused on particular questions.

Damari recounted asking students for their opinions on the ramifications of AI, for instance. He stated:

I think AI ends up a little more in the philosophy world than a lot of people expect just because there are a lot of big implications in it. So I do, probably do a lot more kind of thought pieces. 'And what is your reaction to this? And what is the real world consequences if you were to follow through on this?'

In a similar vein, Cora shared an example line of questioning she might use while leading class discussions about AI. She shared:

I do ask them how they feel about it. What do they see are the pros and the cons, as I, as I introduce any technology, there's both. And you know, what are they excited about? What are they concerned about?

Teachers also touted videos on AI topics as a means to initiate these types of class discussions. For example, Gemma told of sharing videos with students before asking follow up discussion prompt questions. She expressed, "This curriculum, they've provided a lot of videos, you know, so I would just, thought I'd show them this video, and then I would, you know, have some questions. And then some discussion after I showed

them the videos.” Cora described three high interest videos she presented to her students in order to spur discussion. She shared, “I showed three different videos, one on self driving cars, Waymo, is that Waymo? One on Neuralink's work on brain implants and one on ChatGPT.”

When sharing about teacher led pedagogies for teaching AI, participants championed the use of current events and societal issues, teacher-driven class discussion, and video examples to illustrate contemporary AI applications.

Student Tools for Instruction

Not only did participants bring up teaching strategies they find helpful for AI instruction, a group of teachers (N=13) also endorsed tools they find invaluable for AI teaching and learning. Given that teachers found it difficult to teach such a technically complex subject to students who often had little computer programming background, they praised no- and low-code software that illustrate key concepts and allow learners to experiment with AI without the need for advanced coding know-how. For example, teachers reported using Google's Teachable Machine and TensorFlow Playground, Code.org's AI for Oceans, and Concord Consortium's Common Online Data Analysis Platform (CODAP). Itzel explained how this type of software tool allows her to provide a variety of activities during lessons and opportunities for students to enjoy themselves while experimenting with AI. She described:

A lot of the no code tools helped me with that, because they get to kind of play with stuff. So we're not just doing all this heavy, heavy talking, we're also doing a lot of playing and keeping things fun.

Another teacher talked about a specific browser-based tool, Jupyter Notebook, which allows small program segments to be embedded into its interactive IDE. This permits teachers to preload syntax-heavy AI programs into documents for students who can then analyze and experiment with the code snippets rather than having to write the demanding programs themselves. Queen explained:

I prepare Jupyter Notebooks [of AI examples] that I share with them. And then they're, they're able to see that there is some code behind, but they don't have to look at the code. And they can just click, and maybe fill in some text boxes.

Tamika shared her desire to have students learn from writing actual programs without having to navigate the challenges of syntax that written programming languages present. Block-based programming environments have provided this opportunity for her beginning students. She shared, "That's why I love block programming. But because we can, we can do all the computational thinking [of AI algorithms] with none of the silly typographical errors that a lot of my students make [when they write programs using text-based programming languages]."

Teachers also pointed to robotics as an effective interactive tool for students to learn AI. This group emphasized robots' ability to learn and improve as a way to illustrate AI concepts to learners. As Orin put it, "They could start learning about how the robot could teach itself, you know, could could could do some machine learning and it, you know, be an AI in that sense." Ruby also felt that robot programs that self-improve have the potential to illuminate AI for students. She said, "Anybody's going to associate AI and machine learning with robotics. Because the more the robot can teach itself, without having to be explicitly programmed in force, make them more useful."

Discussion of tools vital for AI teaching and learning centered on no- and low-code tools as they allow students to write programs without the burden of complex syntax and robotics platforms for their interactivity and ability to illustrate AI concepts.

Class Communication Norms

Several teachers called for AI instruction to incorporate rich student communication and interaction. These participants stressed the value of putting students in positions to have challenging conversations about AI while in a non-threatening classroom space. They argued that these experiences will empower students to continue to raise concerns about AI at potentially crucial moments throughout their lives. Itzel explained:

The point is not to make people uncomfortable but is to talk about uncomfortable situations so that they are prepared when they go out in the world to have those conversations with people who are maybe not in such a safe place.

Participants also disclosed that because many of their students do not instinctively seek out interactions with their peers, teachers incorporate opportunities for students to practice and strengthen their communication skills into class instruction. For example, Winonah described:

I force my nerdy computer science geeky kids into having to talk to their neighbor, which they would prefer to never have to do. And that's why they signed up for computer science in the first place, thank you very much. And so it forces a little bit of interaction that they're naturally disinclined to do.

Similarly, Fatima reported that she has students participate in collaborative group work where each group member has assigned tasks that include corresponding with

one another. She said, “They also have jobs that they have to do in the groups of four. So forcing them to talk to each other and communicate.” Violet also conveyed the value in students working together in settings that encourage them to communicate with one another to achieve a shared goal. She shared, “I think it's very helpful for them to kind of work together in some smaller groups like that so that they're actually talking to one another.” Itzel also spoke of the importance of setting expectations about how students should handle themselves when challenging discussions about AI arise. She recommended the teacher facilitate norm-setting for students where the class helps to develop guidelines for conversations before engaging in these discussions so students know how to handle discomfort should it arise. She explained:

We do talk in the beginning about conversations can become uncomfortable, and that this needs to be a safe place. And, and what kind of the class itself comes up with? What's a code of conduct? How do we handle ourselves when we talk about things that are difficult if we become uncomfortable?

On the whole, the teachers who shared about class communication norms emphasized the importance of developing guidelines for involving all students in collaborative work and maintaining a safe setting for discussions on sensitive issues.

Unplugged AI Activities

The use of unplugged lessons, or activities that do not make use of a computer, to teach AI concepts also emerged as a theme from the teacher interviews. After teachers experienced the unplugged ML activity described previously in the methods section, they shared a multitude of thoughts on unplugged activities that largely fell into

the following three subcategories: 1) benefits of unplugged activities, 2) keys to unplugged success, and 3) challenges of unplugged activities.

Benefits Of Unplugged Activities

Eighteen teachers spoke about benefits that unplugged activities offer AI learners. These teachers pointed out that unplugged activities can be a promising tool to help students challenge their preexisting assumptions about the field. For instance, Violet extolled unplugged activities as a means to introduce new AI concepts and have students see them in a new light. She described, “Having something early on that makes them [students] stop and do a little more thinking than just going with their instinct.” Another participant, Ruby, emphasized that unplugged activities can sometimes allow students to develop deeper comprehension of concepts than merely using AI tools because they allow students to examine how the underlying technology works. She detailed, “I really think they're better than just going straight to the computer for actually understanding the process of how machine learning works, you can't, you can't get the process just by watching the tool work more.” Similarly, Gemma shared her experience that “Some of these unplugged activities really helped them to think more. And think at a higher level than when they are just on the computer doing stuff.”

Participants also enjoyed that unplugged lessons can be more approachable for students that do not have an abundance of computer programming experience. They explained that CS, and activities that require coding in particular, can be intimidating to many students. However, unplugged lessons typically feel more comfortable since they rarely require students to get bogged down in programming minutiae. One teacher, Bella, explained that since students can approach the lesson without the stress that

comes from doubting their technical knowledge, they can focus on gaining an awareness of the big ideas of the lesson. She shared, "It's understanding the concepts first, and there's no pressure on knowing the technology of it." Violet shared a similar sentiment, saying, "In terms of being able to have those discussions without having to write the stuff behind it, it's really important, I think." Axel felt unplugged lessons allow this because students do not have to carry the cognitive load of the computer programming during the activity. He said:

It helps because it's not code. So they're not they're not trying to figure out the code and see how that they're not bogged down with the code. They're, they're really just experiencing what, what is going to happen when they do AI and how it works without any kind of code.

Elliott clarified that coding can carry with it a stigma that makes it intimidating to many students. Learners are prone to doubting their programming abilities before they try it and in the process, sabotage their learning opportunities during lessons that require coding. Unplugged lessons remedy this problem. He explained:

Coding sort of like, sort of similar to math, some students just kind of tense up at the thought of it. And there's this kind of preconceived notion that the student will have that they're not going to be good at it.

A group of teachers also touted unplugged activities' unique ability to be applied to a wide age range of learners. While most CS lessons are primarily appropriate for a small age band, these teachers explained that the same unplugged lesson is often fitting for elementary, middle, and high school students with little adaptation required. One of those teachers, Liam shared:

I feel like it's appropriate for any level of students, I feel like you could do that in, you know, you could do it in a sixth grade class, you could do it with your seniors at the end of the year, you could do it with anybody.

In summary, teachers praised unplugged activities' power to allow learners to reexamine AI technology in a new light. Participants also shared that students are often less intimidated by unplugged activities since they rarely require computer programming experience. Teachers also appreciated that the lessons can frequently be adapted and applied to a wide age range of learners.

Keys To Unplugged Success

Some teachers (N=10) also shared their fundamentals for delivering successful unplugged lessons. These participants felt that it is key to connect the lesson to things that are important to students. For example, Violet suggested that unplugged lessons should tie into "Something that is, that they [students] can immediately identify with." Damari encouraged teachers and unplugged curriculum developers to consider how they can connect the lesson topic to items that learners are interested in. He said "If you can figure out how to take that thing and tie it to something that's real in their life I think it could be very powerful." Bella suggested the significance of determining what is important to students and tailoring the lesson around it. She stressed that lessons should focus on:

Things that affect their [students'] lives. Like how, I don't do much social media, but I know the girls do Tik Tok and stuff. So how, how is AI being used on Tik Tok? And how can we simulate it in an unplugged activity?

Teachers also emphasized that because AI topics can be intimidating to learners, hands-on unplugged lessons should be designed to make the learning feel more approachable. For instance, Keshawn shared, “I think when it's brought into that tangible, digestible level, it also makes the subject a lot less terrifying or overwhelming for them to learn about.” Violet pointed out that unplugged lessons are especially illustrative when they are built to allow students to take an active role in building their understanding rather than merely absorbing information. She said:

Doing it as an unplugged activity is so much better because they're doing it, as opposed to just you know, hearing about it? ‘Okay, so this is what does,’ it doesn't really, you don't really reflect on what that really means to come up with it. But if you go through those steps of that activity, you're actually having to come up with that model, and then start to think about, ‘Okay, what did I miss in that model?’

Similarly, Sebastian underscored that students respond positively to unplugged activities designed to provide an opportunity for engaging with concepts through manipulatives or other tactile objects. He explained:

It's very hands-on something that the students can can like, grab with their hands, and they, they build it themselves, they test it out themselves, like physically, I think anything that has to do with that, in a computer science class is always like very refreshing for students, because they're not stuck in front of a screen.

Additionally, teachers suggested that unplugged activities should be tailored to students and their current understanding. As Keshawn explained, “It has to meet them

where they are, it's, it has to be something that is very simple to understand to begin.” Participants explained that since many students are motivated by and respond well to games, one way to gear unplugged lessons to students’ interests is to gamify them. Nakeisha suggested that “Really gamifying things really gets them [students] intrigued and engaged.” In addition, Ruby reasoned that “[Unplugged activities] that feel like a game to them, they [students] buy in.” Elliott also recommended that unplugged lessons avoid front-loading a lot of terminology at the beginning of the activity but to introduce it after students experience the concepts. He shared, “If you kind of hide the, the content vocabulary before doing the activity, and then bring it back in at the end.” Elliott also underscored that students often appreciate the opportunity to express themselves and that unplugged lessons should take advantage of this. He said, “They [students] like kind of that ownership and opportunity for creativity.”

When speaking about the fundamentals for success when teaching with unplugged activities, teachers pointed to the value of tying the lesson to students’ interests and also to allowing students to build their own meaning through the activity. Teachers also emphasized the advantages of allowing competition, limiting vocabulary, and allowing students to convey creativity during unplugged activities.

Challenges Of Unplugged Activities

Six teachers also expressed challenges they have encountered when delivering unplugged lessons. Teachers commonly lamented the logistics dilemmas that often accompany unplugged lessons as they can require specialized manipulatives, paper handouts, and adequate space to successfully execute the activities. Winonah explained:

I do think I haven't always loved unplugged stuff. And maybe that could be my bias, could be because it's like a pain in the butt to get kids to do them sometimes. And it could be because I just had to print something out. It's like, 'Oh, crap. Now to get to the printer.' You know, all those logistical things that kind of like, take up that extra bit of time.

Another teacher explained that because time is so limited in a school day unplugged lessons must be simple and quick to prepare for. Liam shared, "Setup has to be easy. So in other words, like, as a teacher, it's gotta be very straightforward, you know, I have a 40-minute class period. So you know, so that's not a ton of time." Ruby discussed the challenge of finding adequate physical space to carry out unplugged lessons. She suggested, "My biggest hurdle in using unplugged group activities is dealing with the space concerns."

Teachers also expressed reservations about unplugged activities' direct focus on the topics being taught. Winonah vocalized that unplugged lessons sometimes skirt around the actual concepts they are meant to teach. She said that they can, "Lack a direct application to the topic at hand. Sometimes I feel like the unplugged lessons, they're a little like topic adjacent." Jose expanded on this idea and explained that unplugged activities are inefficient when they often require teachers to prepare students in advance with extra instruction. He observed:

It definitely requires a lot of preinstruction, right? So they understand what the goal is, like you give them a bunch of materials and then without explaining exactly what the goal of all these materials are, then it would be difficult, it can be difficult. So really knowing what their goal is ahead of time and kind of also not

trying to like, a lot of times the unplugged activities seem to be kind of exploratory or investigative.

Additionally, Damari pointed out that not all topics actually benefit from unplugged style lessons. Instead, he argued that students can experiment with actual AI technology and do so very quickly, something not possible in unplugged activities. He explained:

So if they [students] all have Chromebooks or something like that, then it might make more sense to make this as a digital thing where like, they walk through a decision tree and get to run it a bunch of times, because I think one of the powers of computer science is that you get to try lots of things very fast. But you get to try lots of models and see what the answer is and see what happens on the other side.

In sum, when speaking about using unplugged activities in the classroom, teachers referenced challenges including logistics hurdles, time restrictions, finding lessons that directly apply to the topic being taught, and that particular concepts are better illustrated through the use of a computer.

Challenges of Teaching AI

In addition to discussing the importance of AI education, their fit within the curriculum, and pedagogical approaches, teachers also brought up a variety of challenges to teaching AI at the K-12 level. These challenges fell into the following four subcategories: 1) teacher capacity, 2) student ability limits, 3) traditional structures and processes, and 4) limits of AI curriculum and learning tools.

Teacher Capacity

Eighteen participants shared their concern about teachers' limited capacity to teach AI. Teacher capacity concerns largely focused on the current lack of teacher content knowledge. Many of the inservice teachers (N=12) interviewed perceived their own understanding of AI and ML to be shallow. They argued that this makes it difficult for them to visualize how they should or could teach about the technologies. For example, Fatima lamented, "I just don't know what it is, what it looks like, or what the lessons should be." Her feelings were not isolated as Axel also shared, "I think for me, it'd be more on my end, where I just don't know enough about it, to, you know, put it in, to actually implement it into the curriculum." Likewise, Cora admitted, "I don't have all that knowledge myself right now. So saying how I would teach that, I don't really know." This common response that teachers did not feel prepared to plan and carry out effective AI instruction extended to teachers' lack of confidence in helping their students develop AI and ML comprehension. For instance, Jose expressed worry that his limited understanding could cause him to misinform his students. He shared:

I want to make sure I know enough of it to be able to teach it appropriately and make sure I'm not misleading my students. So I'm a little bit worried about just kind of getting a good mastery of the content myself.

Teachers not only discussed how their lack of AI and ML content knowledge would impact their students' understanding but also what they would be able to teach. Simply put, teachers cannot teach what they do not understand themselves. Bella explained that she would need to teach basic concepts, at least initially, so she could

learn them as preparation for more advanced topics. She said, “I think we have to start with fundamentals, because I'm not familiar, so I want to know about those things, too.”

While teachers expressed interest in learning about AI and ML in order to develop their content knowledge, some (N=10) also had reservations about whether they could find the time to learn themselves as well as fit the subject within a crowded curriculum. For instance, Ruby shared her concern about “Not having time on my own and to, to learn everything that I need to learn in order to teach it.” Another teacher, Tamika, told of her interest in learning and teaching more about AI and ML but could not imagine making enough time to learn it herself. She described, “It's just finding the time. So it's not only not having enough time to teach all the concepts in your course, it's for you for your own professional learning sake, to prepare.” When speaking about the challenges to AI instruction in K-12 education, Axel stated, “I think the main issue is time.” Gemma asked, “Where are they gonna get the people to teach these things? I'm serious. That's like, it's like putting the horse before the carriage. You know, we [teachers] are really, really short [on time to learn and teach AI].” Other teachers elaborated on their struggles with finding enough time for learning about and teaching AI within the structure of the school schedule. Penelope shared, “The biggest challenge is just the amount of time and I'm spread so thin.” Likewise, Damari explained that teachers often feel that in order to make room to teach AI, they must remove other topics from their courses. He explained, “Teachers are already at like a zero sum game. So what is this lesson replacing that they're already doing? Because you have to do that. There's no more time.”

On the whole, teachers expressed unease about their capacity to teach AI and ML because they did not feel prepared or confident in their ability to design AI courses of study given their limited content knowledge. Participants also shared concerns that the amount of teacher planning time and course instructional time available will be inadequate for proper AI instruction.

Student Ability Limits

Teachers were not only hesitant about their own knowledge regarding AI and ML but also of their students' capabilities. Fourteen participants raised worries that students' insufficient abilities could hinder their capacity to learn about AI. One of teachers' primary concerns was students' ability to write AI software given their lack of computer programming experience. Tamika observed, "They're [students] pretty much coming in new [to programming]. We don't have a middle school program." This was not isolated. Violet remarked, "I always survey them [incoming students], 'What have you ever done before?' Rarely have they ever done any programming. And so it's all new." Another teacher explained that teachers should not assume that students will have prior programming knowledge. Queen said, "Because our computer science curriculum is elective, we can't expect students to have programming experience before they join." Cora pointed out that students' limited programming experience should be accounted for when considering what to teach each learner about AI. She felt that it would be challenging to require every student to write AI software when they do not have familiarity with coding. She explained, "I probably wouldn't have everybody in the school take an AI computer science course, you know, where they're actually programming because I think that will require a certain skill level."

Teachers also voiced concerns about whether their high school students' backgrounds would be sufficient to comprehend the inner workings of AI algorithms. Uri called into question students' curiosity about the field and whether they have the concentration necessary. He stated, "A lot of machine learning is very procedural work. And, and the, the amount of focus that's needed for that, I think there's just a few students who have the interest and the focus." Liam focused on whether high school students' math experience would be adequate. He stated, "The trouble there is that their [students'] mathematics level is not high enough." Along the same lines, Queen raised the worry that because AI technologies, such as ML classifiers, are built on sophisticated math, teachers cannot help students truly understand how they work if they do not have the necessary math knowledge which most high students do not yet. She explained, "The math part is challenging, because if we want them [students] to understand how a classifier works, sometimes we have to go deep into math that they don't, they have not learned yet.

The discussion about the challenge of student ability limits centered on learners' lack of experience with computer programming and their limited exposure to advanced mathematics.

Traditional Structures and Processes

A number of teachers (N=13) discussed the difficulties that come with the novelty of AI education. These participants emphasized that K-12 education can be slow to change especially when it comes to the changes necessary for AI education to overcome well ingrained tradition and structures. Many of these participants emphasized the influence that one of these structures - the College Board, the

organization that oversees the AP CS A and Principles exams - has an opinion on what is taught in high school CS classes. As many school's CS offerings include an AP CS course, teachers lamented not having enough time and space to incorporate AI into the classes. For example, Tamika simply stated, "There's no way I can fit it [AI] into AP CSA." Similarly, Axel shared, "[Including AI in] AP Computer Science A would be super hard, just because we're pretty tight on everything that we have to learn. So to veer off and do that might be a little bit tough for the time part." Cora shared that because of the value placed on AP courses, the College Board has an outsized role in determining what is taught in high school CS classes based on what they choose to include on the AP CS tests. She explained:

Because of the power that AP has, you know, that they're kind of controlling what students are being taught all over the country. And that's apparently important for I don't know, if students feel it's important, and colleges feel it's important, this whole AP thing. So there's a lot of power there. And if that doesn't change, it's going to be hard to get, like, AI to be considered an important topic covered in school, unless it's some other non-AP computer science course.

Teachers also talked about the difficulty of overcoming the well established traditions at the school level. Gemma raised concerns that school- and district-level decision-makers are often underinformed about the CS field. She stated, "Administrators don't really know what you're doing, I mean, most administrators these days don't even know what computer science is." She went on to explain that administrators' lack of awareness of the domain makes it difficult for them to make AI

instruction a priority. She added, “Unfortunately, sometimes they [administrators] might not even know enough about it [AI education] to value it.”

Other participants stressed that teachers face almost insurmountable tradition when it comes to teaching practices. Keshawn explained the pressure as follows, “There are a lot of teachers that have a way of teaching, and honestly, they will work 10, 15, 20 years, find a way that works and in their mind and in their lives it is how you teach and there's not a different way.” Winonah raised the point that since veteran educators have taught without AI for years, many felt threatened when the technology became more prevalent. That led some educators to resist all forms of the technology in the classroom including opposing AI instruction in their classrooms and schools. This, in turn, penalizes students who miss out on the opportunity to learn with and about the tools. She explained:

The community in which they're teaching is saying, ‘No, no, we need to ban this from our schools period,’ for whatever reason, and then students are caught in this like area of not either getting the information they need to understand how it works, or getting the boundaries on like, these are the ethical ways in which you can use it.

Participants stressed that AI will necessitate drastic changes to educators’ teaching methods. They argued that teachers need to adopt a radically different mindset in order to take advantage of AI’s knack for simplifying or eliminating tedious work in CS courses and elsewhere. In particular, participants suggested that even though it will require great effort, teachers should completely redesign their lessons to use AI tools to

raise the level of complexity and analysis performed by students. For example, Keshawn recommended:

My goal would be to use it in an assistive manner to be, to do what computers already do for us, to take what we can already create and make it something that's a lot more efficient, a lot more fleshed out by, by removing the tedium and the work that does truly generate no extra Bloom's Taxonomy, right? Doesn't build upon their knowledge. All it does is just kind of make them shy away from the work to begin with.

Participants also talked about teacher resistance to this notion of incorporating AI in order to raise the level of student thinking during assignments in CS and other subjects. Teachers pointed out that it is challenging for established educators to view AI as providing opportunities to allow students to dig deeper into problems rather than ruining their long-established teaching methods. One teacher, Mia, shared her thoughts about teachers at her school, saying:

Everyone [teachers] at lunch is talking about ChatGPT. 'Now it's ruining the essay,' or whatever it is, and I'm like, 'That's just dumb. That's not true. And if it's ruining your essay, then your essay prompt is bad. Like, this is not this is not a computer problem. This is a you problem.'

Liam proposed that one way to help teachers defeat these knee-jerk reactions to AI is to help them recognize how previous ground-breaking technologies have been used by students to advance their thinking and abilities. He suggested that a helpful analogy might be what calculators have allowed students to do in math lessons. He said:

Okay, well, you have a calculator. So now, now, you don't need to teach people to memorize the multiplication table anymore. You know, I said that people don't do long division by hand anymore. But you are now free to do other things and more sophisticated things.

While teachers believed that all students need to understand AI issues, teachers also raised the challenges of recruiting students into their CS courses so they could learn about the technology. In particular, teachers reported difficulty helping students to overcome the well ingrained stereotypes about who belongs in technology fields such as AI. Penelope shared that she finds it challenging to help students from underrepresented populations find their way to CS without being overbearing. She said, "We're trying to break stereotypes in computer science. I know that's a goal. But at the same time, I don't want to push the kids too hard in one direction or the other." Similarly, Violet explained the pervasive sentiment that only nerdy students fit in CS. She described, "There's this feeling out there in the world that you have to be like a math nerd in order to do computer science or engineering. We fight against that a lot and trying to get more students involved." She went on to say that "But still there is just, there's some stereotypes you fight against. And so all of those are the things that you're really working to try to combat as you're trying to get more people interested in it."

In summation, teachers laid out that K-12 AI instruction requires breaking from a variety of traditions and structures including the need to teach the topics found on AP tests, ingrained teaching methods, and stereotypes about who should study technology topics.

Limits of AI Curriculum and Learning Tools

Twelve teachers bemoaned the availability and quality of K-12 AI curriculum and learning tools. Participants emphasized that teachers and schools interested in adding AI instruction for the first time are especially reliant on quality curriculum developed by others. Gemma expressed, “That’s the hardest part, you know, for these schools that are starting up is they got is finding good resources.” Unfortunately, many teachers reported a lack of quality instructional materials aimed at K-12 classrooms. Fatima said, “I just haven’t found a lot of lessons on artificial intelligence.” Likewise, Orin pointed to the challenge of finding learning resources for beginners. He shared, “Availability of tools and sort of starter curriculum, I guess, would be the summary of the challenge that I’m seeing.” Meanwhile, Axel voiced his concern about the absence of unplugged lessons geared toward AI topics. He reported, “I like unplugged activities. And like I said, it’s just a matter of having them available to use.”

While teachers could write the lessons and assessments necessary for classroom instruction on AI, they often feel underqualified to do so since they typically have little experience with the topic. Teachers also have limited time, making full curriculum development challenging. Cora explained that curricula developed by third-parties can solve these problems. She said:

Having those kinds of things (curricula) because teachers just don’t have the time to put this together. They’re not experts in this field. And, you know, having those kinds of things available would be really helpful in helping it get adopted in schools.

Participants who were able to discover AI lessons found that many of them were not optimized for K-12 students. Queen found it challenging to locate learning materials geared toward K-12 rather than post-secondary settings. She said, "I'm always looking for activities that is more at their [K-12 students'] level. So that to expose them to it without having crazy assignments, that would be, you know, university level, college level assignments." Orin also painted a picture of AI curricula he has used. He has found the level of most lessons to be either too challenging or not challenging enough to be engaging. He explained, "The actual work was, it was like, right in that, that terrible middle spot where it was way too hard for some kids, it's way too easy for others."

Teachers suggested that because AI curricula are difficult to find, a centralized repository of learning materials would be helpful. This would also allow for teachers to review materials and provide feedback on their quality. Winonah described it as "One-stop shopping for teachers like, 'Here are the options for you.' Maybe something that's vetted by, you know, people, other teachers some way to, like, communicate this in a streamlined way." Participants emphasized the importance of having both experts in the field and fellow teachers inform what they teach in the classroom. For example, one teacher who has yet to teach AI, shared that she would seek the advice of other teachers about what teaching materials they have had success with before she began teaching the subject. Bella shared:

If I were to teach AI, I'd need to find that community where I could have a mentor, I can have resources of, you know, questions to go back to and what, what curriculum is good for this grade? What worked for you? What didn't - it's having that community, that as a teacher, that really benefits my teaching.

Winonah also spoke of her trust for instructional resources recommended by her peers over those that she stumbles upon. She noted:

I've looked for people that I sort of know from other places, and if they're recommending a tool, or like, Hey, here's a good site, then I'll give it a good look. And I'll tend to trust it. So word of mouth, getting through those, you know, through those Facebook groups, reliable, trusted over the years, and then those are those are kind of the big things.

Additionally, Cora explained that while she could certainly develop her own learning materials, those with AI instructional experience as well as AI subject matter expertise could create a more well-rounded and up to date curriculum. She explained:

I did put together a lot of curriculum, so I'm not opposed to that. This field in particular seems like it's changing so quickly. And, you know, I'm not saying that I have a deep understanding of it. So that's why it would be somewhat reliant on, you know, getting some of the experts to put it together, not just experts in AI, but experts in AI and pedagogy, you know, putting it collaboration, that would be good.

In addition to effective curriculum, participants stressed the importance of technology tools that allow learners to experience AI first hand. Teachers reported a variety of frustrations with access to these resources. Axel shared that the computer hardware his students use to write computer programs limit the complexity of the software they can code. He said:

We have Chromebooks and Chromebooks, you can't code on Chromebooks. We have an online IDE that works, but it's not the greatest kind of thing. You know,

you need a dedicated, dedicated hardware to do some really cool stuff and we just don't have that.

Teachers in schools without these hardware constraints raised other concerns about the availability of software tools appropriate for K-12 AI instruction. Uri commented on the expense of tools developed by third parties and the imperfections of freely available alternatives. He observed:

The problem of a lot of these packaged tools, they're expensive. And the free ones, have those sorts of warts in them, take a long time to fix anything. So you know, I don't feel like I have the tools for it at this point.

Participants also explained that students often do not have access to valuable AI tools due to school or district policies restricting their use. For example, Fatima shared “Yeah, the number one challenge is accessibility.” She went on to share an experience where she tried to encourage other teachers to use an AI tool only to find it blocked on the school network. She recounted:

I was so excited to share it (AI tool) with my other math teachers, but then found out that it was blocked at the school. So it was like, all of my hype for it was nothing. And then so I said, ‘Well, you might want to try it at home.’ And I know that they're not going to because it's just not.

Similarly, Keshawn shared, “Our school district actually has completely restricted access to AI websites on their networks.” He explained that because his school prohibits access on the school-provided internet access, only students who have the financial resources for a data plan on their cell phone are able to access AI tools. He

explained, “The students actually currently have no way to use them [AI tools] outside of their cell phones and cell phone service.”

Sebastian explained that his school’s policies do not allow for students to utilize websites and software, including AI, unless they have been certified compliant with student privacy regulations such as the Children's Online Privacy Protection rule (COPPA). He explained, “We can't access anything unless it's vendor approved or COPPA approved, you know, and so that becomes a challenge.” Uri described a similar challenge with instructional tools that require parent or guardian approval for students under 18 years old. His district does not approve the use of any software with these age restrictions in their terms and conditions of use. This has posed difficulties for finding software that supports AI learning. He outlined:

They're like basic things like Google Colab, which lots of high schools use, have one or two Ts and Cs [terms and conditions] that have an 18 year old minimum, and several that have 13 year old, 13 and over, require below 13 to 18, require permission from parents. And that is now an automatic denial from our IT department.

In essence, many teachers lamented that they did not have the time or expertise to write their own AI curriculum and struggled to find high quality existing instructional materials appropriate for their students. In particular, participants pointed to the challenge of finding curricula for beginners, engaging unplugged activities, and instructional materials recommended by teacher peers. Teachers also experienced limited access to learning resources designed to support K-12 AI instruction.

Underpowered and incompatible computing hardware as well as school policies that restrict teacher and student use of AI tools were also commonly shared issues.

CHAPTER 5: DISCUSSION

This dissertation examined 23 9-12th grade U.S. CS teachers' conceptions and understanding of AI and ML technologies, their perceptions of the field in secondary education, their confidence in teaching AI and ML concepts and practices, and the methods they use as well as the obstacles they encounter while teaching the subject. A discussion of the study's findings, implications, suggested areas for further research, limitations, and conclusions follows.

RQ1 Discussion

How do 9-12th grade CS educators conceptualize AI and ML?

The findings show that teachers had a wide range of understanding of AI and ML but mainly demonstrated surface level comprehension. While these surface level understandings were technically correct, teachers only had a limited knowledge of AI and ML's technical aspects. Teacher views suggested a rudimentary grasp of how AI and ML systems are developed, trained, refined, and function. Additionally, few teachers were able to specify forms of AI and ML or provide considerable example applications of the technologies.

It is interesting to note that many teachers felt confident in their comprehension of AI and ML, but that did not necessarily reflect a deep understanding of the technologies. This misplaced confidence in their own knowledge could be problematic as AI expands into K-12 education if teachers who feel prepared do not in fact have adequate understanding. Meanwhile, even though many teachers were self-assured about their handle of AI and ML concepts, a notable number of teachers were apprehensive about their comprehension of the field and their ability to teach the ideas.

This is especially troubling considering that the study's sample, 9-12th grade CS teachers, is likely more tech savvy than the general K-12 teacher population. If many of the study's participants doubted their understanding of the field and capacity to teach it, non-CS instructors would likely have even lower levels of teacher self-efficacy with regards to AI. Regardless of teachers' confidence level in their understanding of and ability to teach AI, efforts are needed to help them strengthen their knowledge of the field. Teachers must have a deeper understanding of the technology in order to implement AI instruction.

Correspondence with Prior Research

The study's findings on teachers' attitudes and conceptions of AI and ML largely concur with the research base. For example, previous research has shown low teacher confidence in their AI and ML knowledge and preparedness to adequately teach the subject (Chiu & Chai, 2020; Lindner and Berges, 2020; Ng, Lee, et al., 2022; Sanusi et al., 2022). In contrast to Vazhayil et al.'s (2019) findings that most teachers could not see AI career opportunities for their students, results from this study matched previous research that has found that one of the primary drivers for AI instruction is to prepare students for future success (Chiu & Chai, 2020; Ottenbreit-Leftwich et al., 2022). The findings were also consistent with previous literature that inservice teachers want PD to help grow their AI content knowledge (Lindner & Berges, 2020) and that educators are frustrated by the limited availability of AI teaching resources and tools (Chiu & Chai, 2020; Vazhayil et al., 2019). Additionally, teachers feel that ethics instruction is important in K-12 AI education (Ottenbreit-Leftwich et al., 2022) and that they do not have adequate time for AI instruction (Chiu & Chai, 2020). However, some of the

findings did not align with previous research. Ayanwale et al.'s (2022) finding that teachers' perceived confidence in teaching AI predicted their intention to teach the subject and that teachers' belief in the relevance of teaching AI predicted their readiness to teach the topic did not align with results from this study. In the current study, participants with low confidence in their AI teaching ability generally expressed a strong desire to learn more about the field and teach the content. Similarly, while participants saw high relevance in teaching AI, it was not associated with teachers' readiness to teach the subject.

Influence of ChatGPT

In this study, the arrival of ChatGPT greatly informed teachers' conceptions of AI. On November 30, 2022, roughly three months before I began conducting teacher interviews, OpenAI launched the generative AI chatbot Chat Generative Pre-trained Transformer or ChatGPT (OpenAI, 2022). The free-to-use tool exploded in popularity, tallying 100 million monthly active users in a mere two months, the quickest consumer technology application in history to do so (Hu, 2023). Users quickly discovered a myriad of ways to harness ChatGPT's ability to rapidly generate text responses based on simple prompts. Many projected that the tool would transform a wide array of fields and education was no exception (Adıgüzel et al., 2023; Lo, 2023). Early in the interviews it became clear that many participants and their students were some of these early adopters of ChatGPT and other similar generative AI tools based on large language models (LLMs).

Given that ChatGPT is the center of discussion within K-12 schools, it is not surprising that teachers often focused on the software when discussing AI. After

frequently seeing students misuse ChatGPT and other LLM-trained tools, teachers emphasized that learning about AI would allow students to better understand how and when to engage with AI tools. When making a case that students should develop their problem-solving skills instead of merely accepting AI-generated suggestions without question, teachers often referenced students' overuse and misapplication of ChatGPT inside and outside of the CS classroom. Participants raised concerns about the ease of students creating and submitting AI-generated work as their own on computer programming and writing assignments, the challenge of detecting this AI-produced content, and potential learning loss resulting from students' dependence on ChatGPT. At the same time, many of the participants spoke about the importance of their students learning to use AI tools, largely ChatGPT, in wide ranging situations inside and outside of CS.

AI's Risks and Dangers

Not surprisingly, since the technology is new and novel, teachers demonstrated an incomplete perspective of AI's impact on society. While some teachers were able to point out several potential harms (e.g., bias, copyright infringement, data collection, privacy violation), most were not able to describe the risks and dangers that AI technologies pose beyond students potentially becoming over-reliant on them. Researchers have begun to conceptualize ethics in AI and the need for students to understand their role in the necessary structural changes. Given that many of the teachers in the study focused on LLMs and generative AI, there is an opportunity to use this as a starting point to engage CS teachers in how LLMs work and inherent biases within them. For example, Bender et al. (2021) analyzed LLMs, the artificial neural

networks trained on massive amounts of human-generated text which underlie most text-based generative AI tools. The authors termed LLMs *stochastic parrots* to illustrate that LLM text predictions are based on statistical patterns and have limited ability to comprehend their own output. This parroting back of text patterns from human writing leads AI tools to take on the prejudices and biases of those humans who wrote the original training text with limited means to understand that it might be generating incorrect statements including racist, sexist, and xenophobic content. Unfortunately, this type of bias in mainstream technology is not isolated but commonplace. Researchers have long chronicled examples of algorithms discriminating against vulnerable populations (Benjamin, 2020; Eubanks, 2018; Noble, 2018; O'Neil, 2016). However, teachers are largely unfamiliar with this work, and it has not significantly informed their conceptions of AI. As such, there is a need to invest considerable effort to build teachers' awareness and sensitivity of AI ethics so they can work to develop their students' ethical sensitivity as well. In recommendations released in 2023, the U.S. Department of Education advised K-12 schools to provide a more well-rounded view of AI, especially providing opportunities that allow students to critique the technology and its impact on their lives and society as a whole. They argued that examining the ethical issues associated with AI would open students' eyes to dangers such as bias, surveillance, privacy invasion, and security threats (Cardona et al., 2023).

Emerging efforts within this area provide a potential mechanism to develop teacher competencies. For example, the Kapor Foundation's recently released guide on responsible AI and tech justice could serve as a benchmark for the knowledge that teachers require. In particular, the guide focuses on the need for educators to create

opportunities for students to 1) reflect on who designs and benefits from AI, 2) consider AI's impacts on society, 3) critique how data are collected and used, 4) dig into human-computer interactions and their consequences, 5) scrutinize their individual AI use, and 6) work to actively reduce injustice, wrongs, and damage caused by AI (White & Scott, 2024). These six areas could help focus teachers' understanding of ethical and justice-oriented issues within AI education.

RQ2 Discussion

How do 9-12th grade CS educators conceptualize the role of AI in their CS instruction?

The findings show that teachers feel that everyone should understand AI and that the best path for cultivating this understanding is to provide AI literacy instruction for all students during grades K-12. They felt strongly that strengthening students' AI literacy would help learners better recognize, assess, and make use of AI technologies. Additionally, teachers argued that providing AI literacy instruction to every student would help develop a citizenry prepared to make informed decisions about the technology and equip a future workforce for success in an environment rapidly adopting AI systems. Teachers' conceptions about AI literacy's place in K-12 education largely align with the preliminary literature in the field that mostly includes recommendations and frameworks for teaching AI. For example, Digital Promise's newly released AI Literacy Framework emphasized helping learners understand, use, and evaluate AI (Mills et al., 2024). Likewise, Hollands and Breazeal (2024) have encouraged schools to work to introduce students to AI and provide them opportunities to develop their AI literacy before they seek to productively apply tools that employ the technology.

AI Literacy

The study's results also suggest that teachers conceive of AI literacy as a less intimidating way for teachers and students to approach AI rather than immediately beginning with technical instruction. In essence, they viewed AI literacy as an entryway to more advanced AI learning. If all students took part in lessons to build their AI literacy, they would be exposed to AI conceptually and could consider whether they had an interest in pursuing further learning to better understand the technology's mechanics. By and large, participants envisioned that students would best learn the technical side of AI in 9-12th grade elective CS courses. In their current form, these high school CS courses generally incorporate programming instruction and teachers viewed the inclusion of AI coding projects as a valuable opportunity for CS students to apply their programming skills. Including technical AI instruction in 9-12th grade CS courses would also create valuable opportunities to discuss phases in the software design and development cycle where bias might be introduced as students are working through those very phases on their own projects. This would be especially valuable for students who elect to enroll in CS courses as their expressed interest in CS suggests that they need to be more deeply aware of biases in case they design the technologies in the future. These results align well with initial recommendations from state departments of education. Digital Promise recently reviewed the formal guidance on AI in K-12 education issued by seven states and found a common call for AI literacy instruction to equip all students for the future (Roschelle et al., 2024). Several states also suggested that CS courses play an integral part in developing students' understanding of the mechanics of AI (Roschelle et al., 2024).

As K-12 AI instruction ramps up, teachers also need to help students develop a critical perspective of the technology rather than having students see AI as a panacea and cure-all to the problems society faces. Technology that is depicted as objective and neutral on the surface often mirrors human bias and even amplifies discrimination and inequality (Benjamin, 2020; Buolamwini, 2023). This bias has been found in search engine (Noble, 2018), facial recognition (Buolamwini & Gebru, 2018), predictive policing (O’Neil, 2016), welfare benefit (Eubanks, 2018), and LLM (Bender et al., 2021) systems. In order to combat the harm caused by AI systems, teachers must invite students to critique these tools and provide opportunities to reflect on sources and types of bias and prejudice. Learners should understand and appreciate the influence that a diverse range of perspectives can have on the AI design process and level of bias in the final product. Teachers should also empower their students to recognize their role in reversing the biases and improving the software’s design.

AI in Content Standards

Teachers also discussed K-12 content standards both as a blessing and as a curse. On one hand, incorporating AI into content standards would ensure that AI topics are taught in K-12 schools. Embedding AI concepts into the standards of wide-ranging school subjects would provide all students with consistent opportunities to learn about the technology situated in the context of a variety of content areas. Standards would also provide clarity about what should be taught about AI at each grade level and allow teachers to present concepts in a uniform and coherent manner. However, it is important to note that in order to add AI concepts to content standards, other topics would need to drop out of the standards. Schools have fixed instructional time and

resources that require existing subject matter to be eliminated in order to make way for new content in the curriculum. Computing creep, or the charge for AI instruction in this case, could jeopardize fundamental educational content which students still need to learn (Shah & Yadav, 2023).

Teachers also had concerns about the challenge of developing standards for a technology that is rapidly evolving. The standards development process varies by state and is typically a time consuming endeavor that requires input from a multitude of subject matter experts and stakeholder groups, often over the course of 12-24 months. Once approved, the content standards are generally in place for 5-10 years before they are reviewed and updated once again (National Center for Research on Evaluation, Standards, and Student Testing, 2016). Many teachers felt that this type of development cycle is simply inadequate to keep up with the pace of change in AI technology and would require teachers to provide outdated instruction on the field. Additionally, districts, schools, classes, teachers, and students have individual needs; can a one-size-fits-all set of standards truly be sufficient for all of their unique situations?

RQ3 Discussion

What, if any, instructional curricula and pedagogies do 9-12th grade CS educators use to bring AI into their CS instruction?

The findings also show that teachers are craving curricula and tools for teaching AI. Many participants lamented the difficulty of finding curriculum and tools that were a good match for their K-12 students, and in fact, multiple asked hopefully whether an outcome of this study would be ready-to-use software or classroom lessons. Teachers reported that they find it challenging to teach AI concepts in a manner that does not

require advanced programming and/or math knowledge. Participants shared small successes using tools such as Code.org's AI for Oceans and Google's Teachable Machine and TensorFlow Playground but also described their frustration with the lack of software designed to plainly illustrate AI principles to novices. Note that following the teacher interviews, multiple providers, including Code.org and aiEDU, have released more extensive curriculum geared toward K-12 learners; however, much work is left to be done in the curriculum development space.

Importance of Curriculum

The findings also indicate that a trustworthy curriculum not only fills schools' obvious need for teaching materials but could also serve as a map of the AI territory, especially for teachers who are new to the space. Reliable curricula can demonstrate what is universally valuable for students to know and suggest to educators the relative importance of different topics based on the breadth and depth at which they are covered. They also signal what types of pedagogies are effective in different situations and with particular subject matter. While a variety of sources provide viable stand-alone lessons, K-12 teachers are often accustomed to a comprehensive curriculum with interconnected lessons that provide an extensive overview of the field. It is valuable to involve teachers in the curriculum development process to inform the final product. K-12 practitioners can serve an important role and help improve and refine curricula by helping developers better understand their learning environments and students' traits, strengths, and needs (Fullan, 2015; Handler, 2010).

The study's results also point to the importance of teacher confidence in AI curricula. Participants shared criteria they look for in potential lessons. Curricula must

be age appropriate for K-12 students, incorporate instructional best practices, and allow for rich, open-ended discussions. Additionally, curricula must have a low floor making it accessible to all teachers and a high ceiling that opens doors for more technical learning opportunities. Teachers also highlighted the importance of a flexible curriculum that allows them to adapt and pivot to what their students need. Finally, participants emphasized that AI technology changes fast and curricula need to keep up with these updates or they quickly become irrelevant in learners' eyes.

Teaching material repositories are a promising trend that allow teachers to rely on both experts in the field and fellow teachers to inform what they teach in the classroom. These collections, made up of curated resources that have been used in K-12 learning environments, provide teachers an array of instructional material options for teaching AI in a variety of settings. The Hubs in Computing at Schools (National Centre for Computing Education, n.d.) effort in the United Kingdom as well as the Colorado Department of Education's (n.d.) CS Resource Bank and Fairleigh Dickinson University's CS Hub (n.d.) in the U.S. could serve as models to allow K-12 CS teachers to find learning materials.

Pedagogical Approaches

Teachers shared an array of pedagogical strategies they found beneficial for AI instruction. Much emphasis was placed on approaches that position students at the center of the learning such as project-based learning, self-paced learning enabled by instructional videos, class-wide and small group discussions, and interactive programming exercises. These findings largely matched observations from prior research as the top instructional strategies were regularly used in the 32 AI education

studies Yue et al. (2022) reviewed in their systematic literature review. However, multiple participants in the current study expressed a strong aversion to using direct instruction for AI lessons; meanwhile, Yue et al. found direct instruction to be the most common instructional approach used in their review.

Implications

The findings of this study have a variety of implications in the K-12 education space. As few schools, districts, or states have formal initiatives in place to provide AI instruction to students, teachers largely need to take action to bring this instruction to their students until more formal programs are in place. As they do, these findings suggest that it is key that educators seek to broaden and deepen their knowledge of the field in order to turn around and teach their students. A primary goal of their instruction should be to create opportunities for students to see AI's potential for good as well as harm. Ethics should be woven into AI instruction to prepare students to challenge AI technologies that cause harm.

Teacher Training

Participants' AI and ML comprehension varied but generally exhibited a shallow understanding of AI and ML, which illustrates a severe need for PD for teachers. Due to the demands on inservice teachers' time, PD providers should seek to design learning opportunities that not only can scale up to meet sizable demand but are also adaptable to the individual needs and current level of understanding of participants. For maximum benefit to teachers and positive impact on student outcomes, this PD should offer ongoing learning opportunities instead of a single isolated session (Desimone, 2011) and should facilitate collaboration among teachers in professional learning communities

(Borko, 2004; Darling-Hammond, 2010). Ongoing PD and learning communities will ensure that teachers have continued support to truly change teacher understanding and build their capacity (Guskey, 2002).

Teacher educators across all disciplines should also consider how they can incorporate AI literacy into coursework for preservice teachers. These learning opportunities could serve to grow future teachers' awareness of AI and also demonstrate how the technology can be embedded into wide-ranging content areas. The study's results indicate that teachers may not fully recognize the dangers and harms of AI. To overcome this gap in understanding, preservice and inservice teacher training should emphasize the extensive biases, prejudices, threats, and risks posed by AI.

AI Instruction in CS Courses

The results suggest that technical concepts about AI should be incorporated into 9-12th grade CS courses. Given the difficulty of recruiting and training new CS teachers, presenting AI in established CS courses rather than creating new, stand-alone classes is a more practical solution. The U.S. faces a shortage of qualified CS teachers (Basu et al., 2021; Shein, 2019) and inadequate K-12 teacher preparation programs limit the supply of incoming CS educators (Computer Science Teachers Association, 2013). Additionally, new CS teachers have been found to face an abundance of hurdles including narrow content knowledge, inadequate PD, lack of peers, and IT challenges (Yadav et al., 2016). These factors suggest to school leaders that delivering technical AI instruction in existing CS courses with experienced teachers to be a more pragmatic model than creating new CS courses dedicated to AI.

The decision to incorporate AI instruction into CS classes is complicated by the already overcrowded curriculum. Teachers frequently mentioned the vast number of topics they must cover and the difficulty of finding time to add AI instruction. Meanwhile, participants also stressed the importance of teaching AI to prepare students for the future. Unplugged activities may provide viable opportunities to incorporate AI into congested CS courses. The lessons can serve to develop students' basic understanding of AI as well as providing opportunities to discuss ethical issues related to AI.

K-12 AI Curricula

This research also indicates a glaring need for AI curricula geared toward K-12 learners. Curriculum writers can help ground teachers in the space by writing comprehensive curriculum that give students a foundational understanding of AI issues including ethics and bias. The curricula should be tied to existing AI literacy frameworks (Long & Magerko, 2020; Mills et al., 2024) and prepare students to not only be better 21st Century citizens but also allow them to consider whether they want to learn more about how AI works. Teachers suggested that students who express interest in developing technical understanding and skills could learn more in elective 9-12th grade CS courses. Teachers in this study also felt that unplugged lessons showed great promise for illustrating complex AI concepts and helping learners create mental models of abstract principles; thus, curriculum should incorporate unplugged resources which would serve as a low floor for all students to learn AI/ML. Finally, teachers stressed the value of learning materials vetted by other K-12 educators. Groups that support CS educators could help fill this need by creating curriculum repositories that not only allow

teachers to search and view the array of available AI curricula but also review the materials and share peer feedback about their experiences using the instructional resources.

Bias in AI Tools

As teachers and students rapidly embrace ChatGPT and other LLM-based AI tools, they must understand the bias ingrained in the software. LLMs have been shown to reflect the inequalities found in their training data and teachers must be ready to handle the issues that come with their problematic output. Offensive and inappropriate AI-generated content can provide a well-prepared teacher valuable opportunities for students to critique the software and engage in conversations about ethical AI design and use. Instructors can help students consider where ChatGPT and other tools inherit their biases and the need to incorporate diverse perspectives in design, development, and testing. (Kasneci et al., 2023; Ray, 2023).

Implications for Future Research

As K-12 AI education initiatives quickly expand, it will be crucial for researchers to conduct empirical studies to inform the efforts. One priority should be to cultivate approaches that broaden participation to advance equity and diversify the voices that participate in AI design. To further this work, future research should consider how to expand on promising approaches that leverage culturally responsive-sustaining CS education to include AI instruction (Akgun and Greenhow, 2021; Davis et al., 2021; Eglash et al., 2013; Madkins et al., 2019). In addition, future studies should examine how AI literacy could develop students' understanding of AI's potential as well as its limitations.

For instruction to become an important component of K-12 learning as these findings suggest, work should be done to align and consolidate AI curricular frameworks (Lao, 2020; Long & Magerko, 2020; Touretzky et al., 2019). This effort would help teachers, administrators, curriculum developers, teacher educators, and standards writers identify and prioritize essential topics for K-12 AI education. Given that schools are creating an increasing number of CS learning opportunities for elementary-aged students, many learners will enter secondary schools equipped with more advanced computational skills that will improve their ability to learn AI (Touretzky & Gardner-McCune, 2022). Future research could examine the optimal time to teach K-12 students technical AI concepts.

Future studies could also explore how different approaches to integrating AI into wide-ranging classes might help students understand and appreciate the role of the technology across different disciplines. Teachers pointed to the benefit of no- and low-code software tools to facilitate instruction in these interdisciplinary and other efforts to introduce AI to beginners. Future studies could also evaluate the effectiveness of these tools designed to support students' understanding of AI.

As researchers and policymakers work to pinpoint what AI knowledge and skills K-12 students should develop, research must also explore assessment techniques for evaluation of students' comprehension and progress towards these learning targets. Reliable and valid assessment measures designed for K-12 learners will allow educators to pinpoint their learners' level of understanding and adapt their instruction appropriately. Research needs to not only consider what AI competencies are essential for learners but also what knowledge teachers need to lead AI instruction in their

classrooms. This work could expand on prior studies (Kim et al., 2021; Yue et al., 2022) that sought to view AI teacher knowledge through the TPACK framework (Koehler & Mishra, 2009).

Limitations

The study has limitations that need to be acknowledged. First, the qualitative nature of study means that the results are not widely generalizable. The transferability of the findings to a different population or setting should be carefully considered based on the fit or similarity of the situation (Lincoln & Guba, 1985). Due to the small sample size, it cannot be presumed that the conceptions of the teachers interviewed are universally held by all 9-12 grade CS teachers across the U.S. without extensive follow-up research. It is also important to consider the potential for social desirability bias in teachers' responses. Given the nature and context of the study, participants could have surmised that positive comments on AI education would be preferred, leading them to over-emphasize their feelings on the importance of the topic.

In addition, the study's sample consisted solely of 9-12th grade CS teachers teaching at U.S. schools. This intentional choice allowed for a more bounded and precise examination of the issue but limited the generalizability to teachers of other aged students, non-CS subject areas, geographic locations, and educational environments such as informal learning spaces. Moreover, the sample was selected through a call for participation with a large group of teachers, a small number of whom opted in. It is possible that the resulting voluntary sample could have introduced bias since the educators who volunteered for the study may have different traits than the broader population (Wallin, 1949). For example, the teachers who opted into the study

might have more interest in, enthusiasm for, experience with, or comfortability adopting AI instruction than non-participants. Also, recruiting from the membership of the CSTA could lead to a sample that is not entirely representative of the greater CS teacher population.

Conclusion

In his 1903 work *Man and Superman*, Irish playwright George Bernard Shaw opined, “He who can, does. He who cannot, teaches.” This platitude lives on over a century later as teachers are often painted as untalented also-rans who are typically unwilling to go the extra mile for their own or their students’ success. For me, this research eviscerated that narrative. In its place, I came to see that there are many inquisitive, sincere teachers who will do anything in their power to benefit their students. Even when teachers are not equipped with experience in a field and self-assurance in their ability to teach the topic they will get out of their comfort zone to learn along with their students for their benefit. Teaching, especially teaching an unfamiliar discipline, has an incredibly high cognitive load. For educators to successfully prepare their students for an AI-saturated world we must lighten this cognitive load and provide teachers with what they need:

- Curriculum and tools fitting for 9-12 AI instruction
- A grounding in the ethical issues so teachers can help students identify AI’s current flaws and mitigate future bias
- PD to help teachers build their content knowledge, bolster their confidence, and summon the courage to teach an unfamiliar topic

- The time and support teachers need for guided experimentation with the technology

REFERENCES

- Abelson, H., Wolber, D., Morelli, R., Gray, J., & Uche, C. (2012, February). Teaching with App Inventor for Android. In *Proceedings of the 43rd ACM Technical Symposium on Computer Science Education* (pp. 681-681).
<https://doi.org/10.1145/2157136.2157437>
- Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, 13(1), 25-35. <https://doi.org/10.1093/cjres/rsz022>
- Adıgüzel, T., Kaya, M. H., & Cansu, F. K. (2023). Revolutionizing education with AI: Exploring the transformative potential of ChatGPT. *Contemporary Educational Technology*. <https://doi.org/10.30935/cedtech/13152>
- AI4ALL. (2020, February 18). AI4ALL. <https://ai-4-all.org/>
- Akgun, S., & Greenhow, C. (2021). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI and Ethics* (pp. 1-10).
<https://doi.org/10.1007/s43681-021-00096-7>
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211.
[https://doi.org/10.1016/0749-5978\(91\)90020-t](https://doi.org/10.1016/0749-5978(91)90020-t)
- Ali, S., Payne, B. H., Williams, R., Park, H. W., & Breazeal, C. (2019, June). Constructionism, ethics, and creativity: Developing primary and middle school artificial intelligence education. In *International Workshop on Education in Artificial Intelligence K-12 (eduai'19)*, 2, 1-4.
<https://www.media.mit.edu/publications/constructionism-ethics-and-creativity/>
- Angwin, J., Larson, J., Kirchner, L., & Mattu, S. (2016, May 23). *Machine bias*. ProPublica.
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Ayanwale, M. A., Sanusi, I. T., Adelana, O. P., Aruleba, K. D., & Oyelere, S. S. (2022). Teachers' readiness and intention to teach artificial intelligence in schools. *Computers and Education: Artificial Intelligence*, 3, 100099.
<https://doi.org/10.1016/j.caeai.2022.100099>
- Baker, T., Smith, L., and Anissa, N. (2019). *Educ-AI-tion rebooted? Exploring the future of artificial intelligence in schools and colleges*. Nesta.
<https://www.nesta.org.uk/report/education-rebooted>
- Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191. <https://doi.org/10.1037/0033-295x.84.2.191>

- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 671-732. <https://doi.org/10.2139/ssrn.2477899>
- Basu, S., Rutstein, D., & Tate, C. (2021). Building teacher capacity in K–12 computer science by promoting formative assessment literacy. National Comprehensive Center. <https://files.eric.ed.gov/fulltext/ED615714.pdf>
- Bawden, D. (2008). Origins and concepts of digital literacy. *Digital Literacies: Concepts, Policies and Practices*, 30(2008), 17-32.
- Beauchene, V., de Bellefonds, N., Duranton, S., & Mills, S. (2023, June). *AI at work: What people are saying*. Boston Consulting Group. <https://www.bcg.com/publications/2023/what-people-are-saying-about-ai-at-work>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big? 🦜. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623). <https://doi.org/10.1145/3442188.3445922>
- Benjamin, R. (2020). *Race after technology: Abolitionist tools for the new jim code*. Polity Press.
- Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J.M., & Eckersley, P. (2020, January). Explainable machine learning in deployment. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 648-657). <https://doi.org/10.1145/3351095.3375624>
- Blueprint for an AI bill of rights*. (2022, October 4). The White House. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Advances in Neural Information Processing Systems*, 29. https://papers.nips.cc/paper_files/paper/2016/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html
- Borenstein, J., & Howard, A. (2021). Emerging challenges in AI and the need for AI ethics education. *AI and Ethics*, 1, 61-65. <https://doi.org/10.1007/s43681-020-00002-7>
- Borko, H. (2004). Professional development and teacher learning: Mapping the terrain. *Educational Researcher*, 33(8), 3-15. <https://doi.org/10.3102/0013189x033008003>
- Bourke, B. (2014). Positionality: Reflecting on the research process. *The Qualitative Report*, 19(33), 1-9. <https://doi.org/10.46743/2160-3715/2014.1026>

- Bradley, C., Wingfield, R., & Metzger, M. (2020, April). National artificial intelligence strategies and human rights: A review. *London & Stanford*.
https://www.gp-digital.org/wp-content/uploads/2020/04/National-Artificial-Intelligence-Strategies-and-Human-Rights%E2%80%94A-Review_.pdf
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101.
<https://doi.org/10.1191/1478088706qp063oa>
- Braun, V., & Clarke, V. (2012). *Thematic analysis*. American Psychological Association.
- Brehm, L., Guenzel, H., Hinz, O., Humpe, A., & Martius, H. (2019, April). Collaborative learning with COZMO to teach programming in SCRATCH and Python. In *2019 IEEE Global Engineering Education Conference (EDUCON)* (pp. 448-452). IEEE.
<https://doi.org/10.1109/educon.2019.8725037>
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G., Steinhardt, J., Flynn C., hÉigearthaigh, S., Beard, S., Belfield, H., Farquhar, S., ... & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*. <https://arxiv.org/abs/1802.07228>
- Brynjolfsson, E. (2022). The Turing trap: The promise & peril of human-like artificial intelligence. *Daedalus*, 151(2), 272-287. https://doi.org/10.1162/daed_a_01915
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018, May). What can machines learn and what does it mean for occupations and the economy? In *AEA papers and proceedings*, 108, 43-47. <https://doi.org/10.1257/pandp.20181019>
- Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018). Notes from the AI frontier: Modeling the impact of AI on the world economy. *McKinsey Global Institute*, 4.
<https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning>
- Buolamwini, J. 2023. *Unmasking AI: My mission to protect what is human in a world of machines*. Random House.
- Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency* (pp. 77-91). PMLR.
<https://proceedings.mlr.press/v81/buolamwini18a.html>
- Bureau of Labor Statistics, U.S. Department of Labor. (2023, February 6). Computer and information research scientists. In *Occupational Outlook Handbook*.
<https://www.bls.gov/ooh/computer-and-information-technology/computer-and-information-research-scientists.htm>

- Burgsteiner, H., Kandlhofer, M., & Steinbauer, G. (2016). IRobot: Teaching the basics of artificial intelligence in high schools. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1). <https://doi.org/10.1609/aaai.v30i1.9864>
- Cardona, M. A., Rodríguez, R. J., & Ishmael, K. (2023, May). Artificial intelligence and the future of teaching and learning: Insights and recommendations. U.S. Department of Education, Office of Educational Technology. <https://tech.ed.gov/files/2023/05/ai-future-of-teaching-and-learning-report.pdf>
- Carney, M., Webster, B., Alvarado, I., Phillips, K., Howell, N., Griffith, J., Jongejan, J., Pitaru, A., & Chen, A. (2020, April). 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* (pp. 1-8). <https://doi.org/10.1145/3334480.3382839>
- Castro, D., & New, J. (2016). The promise of artificial intelligence. *Center for Data Innovation*, 115(10), 32-35. <https://www2.datainnovation.org/2016-promise-of-ai.pdf>
- Caulfield, H. J., & Maloney, W. T. (1969). Improved discrimination in optical character recognition. *Applied Optics*, 8(11), 2354-2356. <https://doi.org/10.1364/ao.8.002354>
- Chai, C. S., Lin, P.-Y., Jong, M. S.-Y., Dai, Y., Chiu, T. K. F., & Qin, J. (2021). Perceptions of and behavioral intentions towards learning artificial intelligence in primary school students. *Educational Technology & Society*, 24(3), 89-101. <https://www.jstor.org/stable/27032858>
- Charmaz, K. (2014). *Constructing grounded theory*. Sage Publications.
- Chiu, T. K., & Chai, C. S. (2020). Sustainable curriculum planning for artificial intelligence education: A self-determination theory perspective. *Sustainability*, 12(14), 5568. <https://doi.org/10.3390/su12145568>
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 153-163. <https://doi.org/10.1089/big.2016.0047>
- Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. *Wash. L. Rev.*, 89, 1. <https://digitalcommons.law.uw.edu/wlr/vol89/iss1/2/>
- Clarke, V., Braun, V., & Hayfield, N. (2015). Thematic analysis. *Qualitative psychology: A practical guide to research methods*, 3, 222-248. Sage Publications.
- Code.org (2023). *AI for Oceans*. <https://code.org/oceans>

- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1), 37-46.
- Colorado Department of Education. (n.d.). *Computer Science Resource Bank*.
<https://www.cde.state.co.us/computerscience/computer-science-resource-bank>
- Computer Science Teachers Association. (2013). *Bugs in the system: Computer science teacher certification in the U.S.* Association for Computing Machinery.
- Computer Science Teachers Association (2017). *CSTA K-12 computer science standards*. <http://www.csteachers.org/standards>
- Creswell, J. W., & Poth, C. N. (2016). *Qualitative inquiry and research design: Choosing among five approaches*. Sage Publications.
- Curtis, C., Gillespie, N., & Lockey, S. (2022). AI-deploying organizations are key to addressing 'perfect storm' of AI risks. *AI and Ethics*, 1-9.
<https://doi.org/10.1007/s43681-022-00163-7>
- Darling-Hammond, L. (2010). *The flat world and education: How America's commitment to equity will determine our future*. Teachers College Press.
- Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, 6(2), 94.
<https://doi.org/10.7861/futurehosp.6-2-94>
- Davis, K., White, S. V., Dinah, B. C., & Scott, A. (2021). Culturally Responsive-Sustaining Computer Science Education: A Framework. *Kapor Center*.
https://www.kaporcenter.org/wp-content/uploads/2021/07/KC21004_ECS-Frame-work-Report_final.pdf
- Desimone, L. M. (2011). A primer on effective professional development. *Phi Delta Kappan*, 92(6), 68-71. <https://doi.org/10.1177/003172171109200616>
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*. <https://arxiv.org/abs/1810.04805>
- Di Battista, A., Grayling, S., Hasselaar, E., Leopold, T., Li, R., Rayner, M., & Zahidi, S., (2023). Future of jobs report 2023. World Economic Forum, Geneva, Switzerland.
https://www3.weforum.org/docs/WEF_Future_of_Jobs_2023.pdf
- DiPaola, D., Payne, B. H., & Breazeal, C. (2022). Preparing children to be conscientious consumers and designers of AI technologies. In *Computational Thinking Education in K-12: Artificial Intelligence Literacy and Physical Computing*. MIT Press. <https://doi.org/10.7551/mitpress/13375.003.0014>

- Doshi-Velez, F., Kortz, M., Budish, R., Bavitz, C., Gershman, S., O'Brien, D., Scott, K., Schieber, S., Waldo, J., Weinberger, D., Weller, A., & Wood, A. (2017). Accountability of AI under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*. <https://arxiv.org/abs/1711.01134>
- Druga, S. (2018). *Growing up with AI: Cognimates: from coding to teaching machines* [Doctoral dissertation, Massachusetts Institute of Technology]. <https://www.media.mit.edu/publications/growing-up-with-ai/>
- Eglash, R., Gilbert, J. E., & Foster, E. (2013). Toward culturally responsive computing education. *Communications of the ACM*, 56(7), 33-36. <https://doi.org/10.1145/2483852.2483864>
- Eguchi, A. (2021, February). AI-robotics and AI literacy. In *Educational Robotics International Conference* (pp. 75-85). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-77022-8_7
- Ehsan, U., Wintersberger, P., Liao, Q. V., Mara, M., Streit, M., Wachter, S., Riener, A., & Riedl, M. O. (2021). Operationalizing human-centered perspectives in explainable AI. *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3411763.3441342>
- Elliott, A. (2019). *The culture of AI: Everyday life and the digital revolution*. Routledge. <https://doi.org/10.4324/9781315387185>
- England, K. V. (1994). Getting personal: Reflexivity, positionality, and feminist research. *The Professional Geographer*, 46(1), 80-89. <https://doi.org/10.1111/j.0033-0124.1994.00080.x>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- Everson, J., Kivuva, F. M., & Ko, A. J. (2022, February). 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*, 1, 209-215. <https://doi.org/10.1145/3478431.3499395>
- Fast, E., & Horvitz, E. (2017). Long-term trends in the public perception of artificial intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1). <https://doi.org/10.1609/aaai.v31i1.10635>
- Fairleigh Dickinson University CS Hub. (n.d.). *Repository*. <https://www.fducshub.com/repository>
- Fernandes, M. A. (2016). Problem-based learning applied to the artificial intelligence course. *Computer Applications in Engineering Education*, 24(3), 388-399. <https://doi.org/10.1002/cae.21717>

- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation?. *Technological Forecasting and Social Change*, 114, 254-280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Friedman, L., Black, N. B., Walker, E., & Roschelle, J. (2021, November 8). *Safe AI in Education Needs You*. Association of Computing Machinery. <https://cacm.acm.org/blogs/blog-cacm/256657-safe-ai-in-education-needs-you/fulltext>
- Fullan, M. (2015). *The new meaning of educational change*. Teachers College Press.
- Funk, C., Tyson, A., Kennedy, B., & Johnson, C. (2020). Science and scientists held in high esteem across global publics. *Pew Research Center*. https://www.pewresearch.org/science/wp-content/uploads/sites/16/2020/09/PS_2020.09.29_international-science_TOPLINE.pdf
- Furey, H., & Martin, F. (2019). AI education matters: a modular approach to AI ethics education. *AI Matters*, 4(4), 13-15. <https://doi.org/10.1145/3299758.3299764>
- Gardner-McCune, C., Touretzky, D., Martin, F., & Seehorn, D. (2019, February). AI for K-12: Making room for AI in K-12 CS curricula. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education* (pp. 1244-1244). <https://doi.org/10.1145/3287324.3293729>
- Geertz, C. (1973). *Interpretation of Cultures*. Basic Books.
- Gilster, P. (1997). *Digital literacy*. New York: Wiley Computer Pub.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.
- Grbich, C. (2012). *Qualitative Data Analysis: An Introduction (2nd ed.)*. Sage Publications. <https://doi.org/10.4135/9781529799606>
- Gresse von Wangenheim, C., Hauck, J. C., Pacheco, F. S., & Bertonceli Bueno, M. F. (2021). Visual tools for teaching machine learning in K-12: A ten-year systematic mapping. *Education and Information Technologies*, 26(5), 5733-5778. <https://doi.org/10.1007/s10639-021-10570-8>
- Guskey, T. R. (2002). Professional development and teacher change. *Teachers and Teaching*, 8(3), 381-391. <https://doi.org/10.1080/135406002100000512>
- Handler, B. (2010). Teacher as curriculum leader: A consideration of the appropriateness of that role assignment to classroom-based practitioners. *International Journal of Teacher Leadership*, 3(3), 32-42.
- Hanna, P., & Mwale, S. (2017). I'm not with you, yet I am... virtual face-to-face interviews. In Braun, V., Clarke, V., & Gray, D. (Eds.), *Collecting qualitative data:*

A practical guide to textual, media and virtual techniques. Cambridge University Press. <https://doi.org/10.1017/9781107295094.013>

- Heinze, C., Haase, J., & Higgins, H. (2010, July). An action research report from a multi-year approach to teaching artificial intelligence at the K-6 level. In *Proceedings of the AAAI Conference on Artificial Intelligence 24*(3), 1890-1895. <https://doi.org/10.1609/aaai.v24i3.18830>
- Henry, J., Hernalesteen, A., & Collard, A. S. (2021). Teaching artificial intelligence to K-12 through a role-playing game questioning the intelligence concept. *KI-Künstliche Intelligenz*, 35(2), 171-179. <https://doi.org/10.1007/s13218-021-00733-7>
- Holder, C., Khurana, V., & Watts, M. (2018). Artificial intelligence: Public perception, attitude and trust. *Bristows LLP*. <https://www.bristows.com/app/uploads/2019/06/Artificial-Intelligence-Public-Perception-Attitude-and-Trust.pdf>
- Hollands, F., & Breazeal, C. (2024) Establishing AI literacy before adopting AI. *The Science Teacher*, 91(2), 35-42. <https://doi.org/10.1080/00368555.2024.2308316>
- Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base – analyst note. *Reuters*. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Jacobs, M., Pradier, M. F., McCoy, T. H., Perlis, R. H., Doshi-Velez, F., & Gajos, K. Z. (2021). How machine-learning recommendations influence clinician treatment selections: The example of antidepressant selection. *Translational Psychiatry*, 11(1), 1-9. <https://doi.org/10.1038/s41398-021-01224-x>
- Johnson, K., Pasquale, F., & Chapman, J. (2019). Artificial intelligence, machine learning, and bias in finance: toward responsible innovation. *Fordham L. Rev.*, 88, 499.
- K-12 Computer Science Framework Steering Committee. (2016). *K-12 computer science framework*. ACM. <https://k12cs.org/wp-content/uploads/2016/09/K%E2%80%9312-Computer-Science-Framework.pdf>
- Kandlhofer, M., Steinbauer, G., Hirschmugl-Gaisch, S., & Huber, P. (2016). Artificial intelligence and computer science in education: From kindergarten to university. In *IEEE Frontiers in Education Conference* (pp. 1–9). <https://doi.org/10.1109/FIE.2016.7757570>
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T.,

- Stadler, M., Weller, J., Kuhn, J., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kim, S., Jang, Y., Choi, S., Kim, W., Jung, H., Kim, S., & Kim, H. (2021). Analyzing teacher competency with TPACK for K-12 AI education. *KI-Künstliche Intelligenz*, 35(2), 139-151. <https://doi.org/10.1007/s13218-021-00731-9>
- Kezar, A. (2002). Reconstructing static images of leadership: An application of positionality theory. *Journal of Leadership Studies*, 8(3), 94-109. <https://doi.org/10.1177/107179190200800308>
- Koehler, M., & Mishra, P. (2009). What is technological pedagogical content knowledge (TPACK)? *Contemporary Issues in Technology and Teacher Education*, 9(1), 60-70.
- Koepke, L., Brody, D., & Torres, F. (2023, June 7). *Civil Rights Letter to FTC*. Upturn, Inc. <https://www.upturn.org/static/files/2023-06-07-Civil-Rights-Letter-to-FTC.pdf>
- Kong, S. C., Cheung, W. M. Y., & Zhang, G. (2021). Evaluation of an artificial intelligence literacy course for university students with diverse study backgrounds. *Computers and Education: Artificial Intelligence*, 2, 100026. <https://doi.org/10.1016/j.caeai.2021.100026>
- Ladson-Billings, G. (1995). Toward a theory of culturally relevant pedagogy. *American Educational Research Journal*, 32(3), 465-491. <https://doi.org/10.3102/00028312032003465>
- Lander, E., & Nelson, A. (2021, October 22). *ICYMI: WIRED (opinion): Americans need a bill of rights for an AI-powered world*. The White House. <https://www.whitehouse.gov/ostp/news-updates/2021/10/22/icymi-wired-opinion-americans-need-a-bill-of-rights-for-an-ai-powered-world>
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 159-174.
- Lane, D. (2017). *Machine Learning for Kids*. <https://machinelearningforkids.co.uk/>
- Lao, N. (2020). *Reorienting machine learning education towards tinkerers and ML-engaged citizens* [Doctoral dissertation, Massachusetts Institute of Technology]. <https://dspace.mit.edu/handle/1721.1/129264>
- Lee, I., Ali, S., Zhang, H., DiPaola, D., & Breazeal, C. (2021, March). Developing middle school students' AI literacy. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (pp. 191-197). <https://doi.org/10.1145/3408877.3432513>

- Lee, I., & Perret, B. (2022, June). Preparing high school teachers to integrate AI methods into STEM classrooms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11), 12783-12791. <https://doi.org/10.1609/aaai.v36i11.21557>
- Lee, S., Mott, B., Ottenbriet-Leftwich, A., Scribner, A., Taylor, S., Glazewski, K., Hmelo-Silver, C., & Lester, J. (2020, June). Designing a collaborative game-based learning environment for AI-infused inquiry learning in elementary school classrooms. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education* (pp. 566-566). <https://doi.org/10.1145/3341525.3393981>
- Lichtman, M. (2006). *Qualitative research in education: A user's guide*. Sage Publications.
- Lin, P., & Van Brummelen, J. (2021, May). Engaging teachers to co-design integrated AI curriculum for K-12 classrooms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1-12). <https://doi.org/10.1145/3411764.3445377>
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage Publications.
- Lindner, A., & Berges, M. (2020, October). Can you explain AI to me? Teachers' pre-concepts about artificial intelligence. In *2020 IEEE Frontiers in Education Conference (FIE)* (pp. 1-9). IEEE. <https://doi.org/10.1109/fie44824.2020.9274136>
- Lindner, A., Seegerer, S., & Romeike, R. (2019). Unplugged Activities in the Context of AI. In Pozdniakov, S., Dagienė, V. (Eds.) *Informatics in Schools. New Ideas in School Informatics: 12th International Conference on Informatics in Schools: Situation, Evolution, and Perspectives, ISSEP 2019 Proceedings 11913* (123-135). Springer. https://doi.org/10.1007/978-3-030-33759-9_10
- Littman, M. L., Ajunwa, I., Berger, G., Boutilier, C., Currie, M., Doshi-Velez, F., Hadfield, G., Horowitz, M. C., Isbell, C., Kitano, H., Levy, K., Lyons, T., Mitchell, M., Shah, J., Sloman, S., Vallor, S., & Walsh, T. (2021). Gathering strength, gathering storms: The one hundred year study on artificial intelligence (AI100) 2021 study panel report. *Stanford University*. https://ai100.stanford.edu/sites/g/files/sbiybj18871/files/media/file/AI100Report_MT_10.pdf
- Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences*, 13(4), 410. <https://doi.org/10.3390/educsci13040410>
- Long, D., & Magerko, B. (2020, April). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-16). <https://doi.org/10.1145/3313831.3376727>

- Lowry, S., & Macpherson, G. (1988). A blot on the profession. *British Medical Journal (Clinical Research Ed.)*, 296(6623), 657.
<https://doi.org/10.1136/bmj.296.6623.657>
- Madkins, T. C., Martin, A., Ryoo, J., Scott, K. A., Goode, J., Scott, A., & McAlear, F. (2019, February). Culturally relevant computer science pedagogy: From theory to practice. In *2019 Research on Equity and Sustained Participation in Engineering, Computing, and Technology (RESPECT)* (pp. 1-4). IEEE.
<https://doi.org/10.1109/respect46404.2019.8985773>
- Makridakis, S. (2017). The forthcoming Artificial Intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90, 46-60.
<https://doi.org/10.1016/j.futures.2017.03.006>
- Malterud, K. (2001). Qualitative research: standards, challenges, and guidelines. *The Lancet*, 358(9280), 483-488. [https://doi.org/10.1016/s0140-6736\(01\)05627-6](https://doi.org/10.1016/s0140-6736(01)05627-6)
- Manheim, K., & Kaplan, L. (2019). Artificial intelligence: Risks to privacy and democracy. *Yale JL & Tech.*, 21, 106.
https://yjolt.org/sites/default/files/21_yale_j.l._tech._106_0.pdf
- Marques, L. S., Gresse von Wangenheim, C., & Hauck, J. C. (2020). Teaching machine learning in school: A systematic mapping of the state of the art. *Informatics in Education*, 19(2), 283-321. <https://doi.org/10.15388/infedu.2020.14>
- Mason, J. (2017). *Qualitative researching*. Sage Publications.
- Maxwell, J. A. (2012). *Qualitative research design: An interactive approach*. Sage Publications.
- McCarthy, J. (2007). What is artificial intelligence?
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). A proposal for the Dartmouth summer research project on artificial intelligence.
<http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- McCorduck, P., & Cfe, C. (2004). *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. CRC Press.
<https://doi.org/10.2307/1574702>
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An expanded sourcebook*. Sage Publications.
- Mills, K., Ruiz, P., & Lee, K. (2024, February 21). *Revealing an AI literacy framework for learners and educators*. Digital Promise.
<https://digitalpromise.org/2024/02/21/revealing-an-ai-literacy-framework-for-learners-and-educators/>

- Mitchell, T. M. (1997). *Machine learning* (Vol. 1, No. 9). New York: McGraw-hill.
- Morning Consult/Politico. (2017, April 1). *Morning consult national tracking poll #170401*.
https://www.orfonline.org/wp-content/uploads/2019/09/ORF_Issue_Brief_313_AI_Norms.pdf
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. MIT press.
- Murphy, R. F. (2019). Artificial intelligence applications to support K-12 teachers and teaching. *Rand Corporation*, 10. <https://doi.org/10.7249/pe315>
- National Center for Research on Evaluation, Standards, and Student Testing. (2016). States' content standards revision processes. *Eric.ed.gov*.
<https://files.eric.ed.gov/fulltext/ED588498.pdf>
- National Centre for Computing Education. (n.d.). *Hubs in Computing at Schools*.
<https://teachcomputing.org/hubs>
- National Governors Association Center for Best Practices, Council of Chief State School Officers (2010). Common core state standards. *Corestandards.org*.
<http://www.corestandards.org/read-the-standards/>
- Ng, A. (2016). What artificial intelligence can and can't do right now. *Harvard Business Review*, 9(11).
<https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now>
- Ng, D. T. K., Lee, M., Tan, R. J. Y., Hu, X., Downie, J. S., & Chu, S. K. W. (2022). A review of AI teaching and learning from 2000 to 2020. *Education and Information Technologies*, 1-57. <https://doi.org/10.1007/s10639-022-11491-w>
- Ng, D. T. K., Leung, J. K. L., Chu, K. W. S., & Qiao, M. S. (2021a). AI literacy: Definition, teaching, evaluation and ethical issues. *Proceedings of the Association for Information Science and Technology*, 58(1), 504–509.
<https://doi.org/10.1002/pr2.487>
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021b). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- Ng, D. T. K., Luo, W., Chan, H. M. Y., & Chu, S. K. W. (2022). Using digital story writing as a pedagogy to develop AI literacy among primary students. *Computers and Education: Artificial Intelligence*, 3, 100054.
<https://doi.org/10.1016/j.caeai.2022.100054>
- Ng, D. T. K., Wu, W., Leung, J. K. L., & Chu, S. K. W. (2023, July). Artificial Intelligence (AI) literacy questionnaire with confirmatory factor analysis. In *2023 IEEE*

- International Conference on Advanced Learning Technologies (ICALT)* (pp. 233-235). IEEE. <https://doi.org/10.1109/icalt58122.2023.00074>
- NGSS Lead States (2013). *Next generation science standards: For states, by states*. The National Academies Press.
<https://www.nextgenscience.org/standards/standards>
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
- Norouzi, N., Chaturvedi, S., & Rutledge, M. (2020, April). Lessons learned from teaching machine learning and natural language processing to high school students. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 34(9), 13397-13403. <https://doi.org/10.1609/aaai.v34i09.7063>
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books. <https://doi.org/10.14507/er.v24.2197>
- Omaar, H. (2022, July 27). U.S. AI policy report card. *Center for Data Innovation*.
<https://datainnovation.org/2022/07/ai-policy-report-card/>
- Omiye, J. A., Lester, J. C., Spichak, S., Rotemberg, V., & Daneshjou, R. (2023). Large language models propagate race-based medicine. *NPJ Digital Medicine*, 6(1), 195. <https://doi.org/10.1038/s41746-023-00939-z>
- OpenAI. (2022, November 30). Introducing ChatGPT. OpenAI.
<https://openai.com/blog/chatgpt>
- Orth, T., & Bialik, C. (2023, April 20). *Most Americans think schools should teach students to use AI appropriately, rather than banning it*. YouGov.
<https://today.yougov.com/topics/technology/articles-reports/2023/04/20/most-think-schools-should-teach-about-ai-poll>
- Ottenbreit-Leftwich, A., Glazewski, K., Jeon, M., Jantaraweragul, K., Hmelo-Silver, C. E., Scribner, A., Lee, S., Mott, B., & Lester, J. (2022). Lessons learned for AI education with elementary students and teachers. *International Journal of Artificial Intelligence in Education*, 1-23.
<https://doi.org/10.1007/s40593-022-00304-3>
- Papert, S. A. (1980). *Mindstorms: Children, computers, and powerful ideas*. Basic Books.
- Payne, B. H. (2019). An ethics of artificial intelligence curriculum for middle school students. *MIT Media Lab Personal Robots Group*.
<https://thecenter.mit.edu/wp-content/uploads/2020/07/MIT-AI-Ethics-Education-Curriculum.pdf>

- Qian, Y., & Lehman, J. (2017). Students' misconceptions and other difficulties in introductory programming: A literature review. *ACM Transactions on Computing Education (TOCE)*, 18(1), 1-24. <https://doi.org/10.1145/3077618>
- Quach, K. (2020, July 1). *MIT apologizes, permanently pulls offline huge dataset that taught AI systems to use racist, misogynistic slurs*. The Register. https://www.theregister.com/2020/07/01/mit_dataset_removed
- Rainie, L., & Anderson, J. (2017, February 8). *Code-dependent: Pros and cons of the algorithm age*. Pew Research Center: Internet, Science & Tech. <https://www.pewresearch.org/internet/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/>
- Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., & Kiron, D. (2019). Winning with AI: Pioneers combine strategy, organizational behavior, and technology. *MIT Sloan Management Review*. <https://sloanreview.mit.edu/projects/winning-with-ai/>
- Raub, M. (2018). Bots, bias and big data: artificial intelligence, algorithmic bias and disparate impact liability in hiring practices. *Ark. L. Rev.*, 71, 529. <https://scholarworks.uark.edu/cgi/viewcontent.cgi?article=1052&context=alr>
- Ray, P. P. (2023). ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things and Cyber-Physical Systems*, 3, 121-154. <https://doi.org/10.1016/j.iotcps.2023.04.003>
- Resnick, M., Maloney, J., Monroy-Hernández, A., Rusk, N., Eastmond, E., Brennan, K., Millner, A., Rosenbaum, E., Silver, J., Silverman, B., & Kafai, Y. (2009). Scratch: programming for all. *Communications of the ACM*, 52(11), 60-67. <https://doi.org/10.1145/1592761.1592779>
- Rizvi, S., Waite, J., & Sentance, S. (2023). Artificial Intelligence teaching and learning in K-12 from 2019 to 2022: A systematic literature review. *Computers and Education: Artificial Intelligence*, 100145. <https://doi.org/10.1016/j.caeai.2023.100145>
- Roberts, J. K., Pavlakis, A. E., & Richards, M. P. (2021). It's more complicated than it seems: Virtual qualitative research in the COVID-19 era. *International Journal of Qualitative Methods*, 20, 16094069211002959. <https://doi.org/10.1177/16094069211002959>
- Roschelle, J., Lester, J., & Fusco, J. (2020, February 24). *AI and the future of learning: Expert panel report*. Digital Promise. <https://doi.org/10.51388/20.500.12265/106>
- Roschelle, J., Fusco, J., & Ruiz, P. (2024). *Review of Guidance from Seven States on AI in Education*. Digital Promise. <https://doi.org/10.51388/20.500.12265/204>

- Rozado, D. (2023). Danger in the machine: The perils of political and demographic biases embedded in AI systems. *Manhattan Institute*.
<https://media4.manhattan-institute.org/sites/default/files/the-perils-of-political-and-demographic-biases-embedded-in-ai-systems.pdf>
- Rubin, H. J., & Rubin, I. S. (2011). *Qualitative interviewing: The art of hearing data*. Sage Publications. <https://doi.org/10.4135/9781452226651>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach, global edition* (4th ed.). Pearson Education.
- Sabuncuoglu, A. (2020, June). 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* (pp. 96-102).
<https://doi.org/10.1145/3341525.3387364>
- Sadler, P. M., Sonnert, G., Coyle, H. P., Cook-Smith, N., & Miller, J. L. (2013). The influence of teachers' knowledge on student learning in middle school physical science classrooms. *American Educational Research Journal*, 50(5), 1020-1049.
<https://doi.org/10.3102/0002831213477680>
- Saldaña, J. (2021). *The coding manual for qualitative researchers*. Sage Publications.
- Samuel, A. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3), 210-229.
<https://doi.org/10.1147/rd.33.0210>
- Sanusi, I. T., Oyelere, S. S., Agbo, F. J., & Suhonen, J. (2021, October). Survey of resources for introducing machine learning in K-12 context. In *2021 IEEE Frontiers in Education Conference (FIE)* (pp. 1-9). IEEE.
<https://doi.org/10.1109/fie49875.2021.9637393>
- Sanusi, I. T., Oyelere, S. S., & Omidiora, J. O. (2022). Exploring teachers' preconceptions of teaching machine learning in high school: A preliminary insight from Africa. *Computers and Education Open*, 3, 100072.
<https://doi.org/10.1016/j.caeo.2021.100072>
- Sarker, I. H. (2021). Machine learning: Algorithms, real-world applications and research directions. *SN Computer Science*, 2(3), 160.
<https://doi.org/10.1007/s42979-021-00592-x>
- Sarker, I. H., Furhad, M. H., & Nowrozy, R. (2021). Ai-driven cybersecurity: An overview, security intelligence modeling and research directions. *SN Computer Science*, 2(3), 1-18. <https://doi.org/10.1007/s42979-021-00557-0>
- Scheidt, A., & Pulver, T. (2019). Any-cubes: A children's toy for learning AI: Enhanced play with deep learning and MQTT. In *Proceedings of Mensch und Computer 2019* (pp. 893-895). <https://doi.org/10.1145/3340764.3345375>

- Shah, N., & Yadav, A. (2023). Racial justice amidst the dangers of computing creep: A dialogue. *TechTrends*, 67(3), 467-474.
<https://doi.org/10.1007/s11528-023-00835-z>
- Shaw, G. B. (1903). *Man and Superman*.
- Shein, E. (2019). The CS teacher shortage. *Communications of the ACM*, 62(10), 17-18. <https://doi.org/10.1145/3355375>
- Shenton, A. K. (2004). Strategies for ensuring trustworthiness in qualitative research projects. *Education for Information*, 22(2), 63-75.
<https://doi.org/10.3233/efi-2004-22201>
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.
<https://doi.org/10.1038/nature16961>
- Slonim, N., Bilu, Y., Alzate, C., Bar-Haim, R., Bogin, B., Bonin, F., Choshen, L., Cohen-Karlik, E., Dankin, L., Edelstein, L., Ein-Dor, L., Friedman-Melamed, R., Gavron, A., Gera, A., Gleize, M., Gretz, S., Gutfreund, D., Halfon, A., Hershovich, D., ... Aharonov, R. (2021). An autonomous debating system. *Nature*, 591(7850), 379–384. <https://doi.org/10.1038/s41586-021-03215-w>
- Smilkov, D., Carter, S., Sculley, D., Viégas, F. B., & Wattenberg, M. (2017). Direct-manipulation visualization of deep networks. *arXiv preprint arXiv:1708.03788*. <https://arxiv.org/abs/1708.03788>
- Steinbauer, G., Kandlhofer, M., Chklovski, T., Heintz, F., & Koenig, S. (2021). A differentiated discussion about AI education K-12. *KI-Künstliche Intelligenz*, 35(2), 131-137. <https://doi.org/10.1007/s13218-021-00724-8>
- Stolee, K. T., & Fristoe, T. (2011, March). Expressing computer science concepts through Kodu game lab. In *Proceedings of the 42nd ACM Technical Symposium on Computer Science Education* (pp. 99-104).
<https://doi.org/10.1145/1953163.1953197>
- Strauss, A. L. (1987). *Qualitative analysis for social scientists*. Cambridge University Press.
- Stringer, E. T., & Aragón, A. O. (2020). *Action research*. Sage Publications.
- Sulmont, E., Patitsas, E., & Cooperstock, J. R. (2019, February). Can you teach me to machine learn? In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education* (pp. 948-954).
<https://doi.org/10.1145/3287324.3287392>

- Suresh, H., & Guttag, J. (2021). A framework for understanding sources of harm throughout the machine learning life cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization* (pp. 1-9).
<https://doi.org/10.1145/3465416.3483305>
- Susser, D. (2019, January). Invisible influence: Artificial intelligence and the ethics of adaptive choice architectures. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 403-408).
<https://doi.org/10.1145/3306618.3314286>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Swoboda, N., Bekios-Calfa, J., Baumela, L., & de Lope, J. (2011, March). An introduction to AI course with guide robot programming assignments. In *Proceedings of the 42nd ACM Technical Symposium on Computer Science Education* (pp. 231-236). <https://doi.org/10.1145/1953163.1953231>
- Thompson, D., Bell, T., Andreae, P., & Robins, A. (2013, March). The role of teachers in implementing curriculum changes. In *Proceedings of the 44th ACM technical symposium on Computer science education* (pp. 245-250).
<https://doi.org/10.1145/2445196.2445272>
- Tissenbaum, M., Sheldon, J., & Abelson, H. (2019). From computational thinking to computational action. *Communications of the ACM*, 62(3), 34-36.
<https://doi.org/10.1145/3265747>
- Toivonen, T., Jormanainen, I., Kahila, J., Tedre, M., Valtonen, T. & Vartiainen, H. (2020). Co-designing machine learning apps in K-12 with primary school children. In *2020 IEEE: 20th International Conference on Advanced Learning Technologies* (pp. 308–310). IEEE. <https://doi.org/10.1109/icalt49669.2020.00099>
- Touretzky, D., & Gardner-McCune, C. (2018, February). Calypso for Cozmo: Robotic AI for everyone. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education* (pp. 1110-1110).
<https://doi.org/10.1145/3159450.3162200>
- Touretzky, D., & Gardner-McCune, C. (2022). Chapter IX: Artificial Intelligence Thinking in K-12. In Kong, S. C., & Abelson, H. (Eds.), *Computational Thinking in K-12: Artificial Intelligence Literacy and Physical Computing*. MIT Press.
<https://doi.org/10.7551/mitpress/13375.003.0013>
- Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019, July). Envisioning AI for K-12: What should every child know about AI?. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 9795-9799.
<https://doi.org/10.1609/aaai.v33i01.33019795>
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.
<https://doi.org/10.1093/mind/lix.236.433>

- UNESCO (Ed.). (2022). *K-12 AI curricula: A mapping of government-endorsed AI curricula*. <https://unesdoc.unesco.org/ark:/48223/pf0000380602>
- Van Rossum, G. (2007, June). Python programming language. In *USENIX Annual Technical Conference*, 41(1), 1-36.
- Vazhayil, A., Shetty, R., Bhavani, R. R., & Akshay, N. (2019, December). Focusing on teacher education to introduce AI in schools: Perspectives and illustrative findings. In *2019 IEEE Tenth International Conference on Technology for Education (T4E)*, 71-77. IEEE. <https://doi.org/10.1109/t4e.2019.00021>
- VB Staff. (2021, December 6). *Report: AI investments see largest year-over-year growth in 20 years*. VentureBeat. <https://venturebeat.com/ai/report-ai-investments-see-largest-year-over-year-growth-in-20-years/>
- Wallin, P. (1949). Volunteer subjects as a source of sampling bias. *American Journal of Sociology*, 54(6), 539-544. <https://doi.org/10.1086/220418>
- Wang, N., Tonko, P., Ragav, N., Chungyoun, M., & Plucker, J. (2022). A perspective on K-12 AI education. *Technology & Innovation*. <https://doi.org/10.21300/23.1.2023.2>
- Wang, P. (1995). *Non-axiomatic reasoning system: Exploring the essence of intelligence*. [Doctoral dissertation, Indiana University]. <https://cis.temple.edu/~wangp/Publication/thesis.pdf>
- Wang, P. (2019). On defining artificial intelligence. *Journal of Artificial General Intelligence*, 10(2), 1-37. <https://doi.org/10.2478/jagi-2019-0002>
- West, D. M., & Allen, J. R. (2018, April 24). Report: How artificial intelligence is transforming the world. The Brookings Institution. <https://www.brookings.edu/research/how-artificial-intelligence-is-transforming-the-world/>
- White, S. V. , & Scott, A. (2024, January 18). *Responsible AI & tech justice: A guide for K-12 education*. Kapor Foundation. <https://kaporfoundation.org/wp-content/uploads/2024/01/Responsible-AI-Guide-Kapor-Foundation.pdf>
- Wiggins, G. P., & McTighe, J. (2005). *Understanding by design* (2nd ed.). Pearson.
- Williams, R., Park, H. W., & Breazeal, C. (2019, May). A is for artificial intelligence: the impact of artificial intelligence activities on young children's perceptions of robots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-11. <https://doi.org/10.1145/3290605.3300677>

- Yadav, A., Gretter, S., Hambrusch, S., & Sands, P. (2016). Expanding computer science education in schools: understanding teacher experiences and challenges. *Computer science education*, 26(4), 235-254. <https://doi.org/10.1080/08993408.2016.1257418>
- Yadav, A., & Heath, M. K. (2022). Breaking the code: Confronting racism in computer science through community, criticality, and citizenship. *TechTrends*, 1-9. <https://doi.org/10.1007/s11528-022-00734-9>
- Yau, K. W., Chai, C. S., Chiu, T. K., Meng, H., King, I., & Yam, Y. (2023). A phenomenographic approach on teacher conceptions of teaching Artificial Intelligence (AI) in K-12 schools. *Education and Information Technologies*, 28(1), 1041-1064. <https://doi.org/10.1007/s10639-022-11161-x>
- Yue, M., Jong, M. S. Y., & Dai, Y. (2022). Pedagogical design of K-12 artificial intelligence education: A systematic review. *Sustainability*, 14(23), 15620. <https://doi.org/10.3390/su142315620>
- Yurtsever, E., Lambert, J., Carballo, A., & Takeda, K. (2020). A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8, 58443-58469. <https://doi.org/10.1109/access.2020.2983149>
- Zahidi, S., Ratcheva, V., Hingel, G., & Brown, S. (2020). The future of jobs report 2020. World Economic Forum. <https://www.weforum.org/reports/the-future-of-jobs-report-2020>
- Završnik, A. (2020). Criminal justice, artificial intelligence systems, and human rights. *ERA Forum*, 20(4), 567–583. <https://doi.org/10.1007/s12027-020-00602-0>
- Zhang, B., & Dafoe, A. (2019). Artificial intelligence: American attitudes and trends. Available at SSRN 3312874. <https://doi.org/10.2139/ssrn.3312874>
- Zhang, D., Maslej, N., Brynjolfsson, E., Etchemendy, J., Lyons, T., Manyika, J., Ngo, H., Niebles, J. C., Sellitto, M., Sakhaee, E., Shoham, Y., Clark, J., & Perrault, R. (2022). The AI index 2022 annual report, AI Index Steering Committee. <https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report-Master.pdf>. <https://doi.org/10.1108/978-1-80262-633-920221008>
- Zhang, H., Lee, I., Ali, S., DiPaola, D., Cheng, Y., & Breazeal, C. (2022). Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. *International Journal of Artificial Intelligence in Education*, 1-35. <https://doi.org/10.1007/s40593-022-00293-3>
- Zhou, X., Van Brummelen, J., & Lin, P. (2020). Designing AI learning experiences for K-12: Emerging works, future opportunities and a design framework. *arXiv preprint arXiv:2009.10228*. <https://arxiv.org/abs/2009.10228>

Zou, L., & Khern-am-nuai, W. (2022). AI and housing discrimination: The case of mortgage applications. *AI and Ethics*, 1-11.
<https://doi.org/10.1007/s43681-022-00234-9>

Zweben, S., & Bizot, B (2022). 2022 CRA Taulbee survey. *Computing Research News*, 34(5), 2-82. <https://cra.org/resources/taulbee-survey/>

APPENDIX A: RECRUITMENT MESSAGE

Subject: Research Study: Seeking US High School CS Teachers

We are interested in high school computer science teachers to participate in a study to better understand how they see the role of AI in CS classrooms.

Participation in the study requires taking a 5- to 10-minute background survey and engaging in a 60- to 90-minute interview. You will be compensated with a \$100 Amazon gift card if you are selected and complete the study.

If you would like to volunteer, please indicate your interest by filling out the background survey ([link to questionnaire here](#)).

If you have questions or concerns, please let me know ([email address listed here](#)).
Thank you.

Zac Opps
Graduate Student
Michigan State University

APPENDIX B: PARTICIPANT SCREENING SURVEY

Potential participants filled out the following Qualtrics questionnaire to express their interest in participating in the study:

Screen 1

This survey is part of a research project about high school Computer Science (CS) teachers' thoughts on artificial intelligence and machine learning. Thank you for expressing your interest in participating in the study. The following questions will be used to help the researcher select a variety of interviewees. Your name and contact information will only be used to inform you whether you have been selected to participate and will not be shared.

If you have questions or concerns, please contact Zac Opps (oppszach@msu.edu). Thank you.

Screen 2

How long have you been teaching CS content?

- 0-1 year
- 2-4 years
- 5-9 years
- 10+ years

How did you prepare to teach CS content (select all that apply)?

- CS Education degree
- CS (not education degree)
- Professional development
- Self-taught through online tutorials, etc.
- Other

What CS courses do you teach?

What grades are your students in the CS courses?

What other (non-CS) courses do you teach, if any?

Do you integrate CS into these other courses, and if so, please explain how you do this?

What grades are your students in the non-CS courses?

Screen 3

Your name

Your email address

Preferred pronouns

She / her / hers

He / him / his

They / them / theirs

Other

Race / ethnicity

American Indian or Alaska Native

Asian

Black or African American

Hispanic or Latino

Native Hawaiian or other Pacific Islander

White

Other

What is the estimated student enrollment of the school you teach in?

0 - 249

250 - 1000

1000 - 2000

2000+

What is the setting of the school you teach in?

Rural

Suburban

Urban

Do you teach in a Title I school?

Yes

No

Unsure

APPENDIX C: CONSENT FORM

RESEARCH PARTICIPANT INFORMATION AND CONSENT FORM

You are being asked to participate in a study carried out by Michigan State University. Researchers working on this project are required to present a consent form that gives you the option to participate or not. This form details the purpose of the study and the potential risks and benefits of participating. You are free to ask the researchers any questions about the program and opt out of the study at any time now or in the future, even after providing consent.

Study Title: Teacher Perceptions of Artificial Intelligence

Researchers: Zac Opps & Aman Yadav

Institution: Counseling, Educational Psychology and Special Education
College of Education, Michigan State University
620 Farm Lane, Room 513A, East Lansing, MI 48824

Phone: 307-250-5902

Email: oppszach@msu.edu

1. PURPOSE OF RESEARCH:

The purpose of this study is to explore teachers' perceptions of artificial intelligence and machine learning.

2. WHAT YOU WILL DO:

Consenting to be part of the research project means that you agree to participate in a 60- to 90-minute video call interview which will be recorded for future transcription and analysis. If you consent to participate, you may opt-out at any time.

3. POTENTIAL BENEFITS:

Being a part of this project will assist you in understanding your own perspectives on artificial intelligence and machine learning education. This may help you better consider how these technologies impact your students and matters relevant to teaching about the technologies.

4. POTENTIAL RISKS:

All your individual and personal information will be anonymized and kept confidential at all stages of this research. No potential risks are foreseen as a result.

5. PRIVACY AND CONFIDENTIALITY:

Data collected from this research project will be kept confidential at all stages of research until aggregated and anonymized for presentation in academic publications and conferences, at which time pseudonyms will be used for individual names, locations, and associated products. Only members of the research team will have access to the raw data. All video, audio, and textual data will be stored on password required hard drives and locked in file cabinets when not being used by the research team.

6. YOUR RIGHTS TO PARTICIPATE, SAY NO, OR WITHDRAW:

Participation in this research project is completely voluntary. All participants have the right to say no and refuse data to be collected on them. In addition, you may withdraw from the project at any time and for any reason. You may choose not to participate at all, or you may refuse to participate in certain procedures or answer certain questions or discontinue your participation at any time without consequence.

7. COMPENSATION FOR BEING IN THE STUDY:

You will be compensated with a \$100 gift card for participating in the interview.

8. CONTACT INFORMATION FOR QUESTIONS AND CONCERNS:

Questions and concerns about participating in this study can be directed to the researchers. Please find their contact information above.

9. INFORMED CONSENT:

You indicate your voluntary agreement to participate in this study by beginning the video call interview.

APPENDIX D: INTERVIEW PROTOCOL

The semi-structured interviews were guided by the following protocol:

1. *Read introduction script* - This study seeks to have conversations with CS teachers to find out what they think about artificial intelligence & machine learning. I want to emphasize that there are no “right” or “wrong” answers to any of the questions we will discuss today. We want to be sure that you are informed about your role and rights in the research process. Did you have a chance to read the consent form (if not, will you please take a few minutes right now to read through it before we get started?). Do you have any questions or reservations about anything you read on the form? The interview will be recorded so we can focus on our conversation rather than notetaking but the video and all of your answers will remain confidential. Our conversation will likely last roughly 60 to 90 minutes. Do you have any questions or concerns before we begin?
2. *Rapport building* - Welcome. Thank you so much for taking the time to talk with me today. As a former CS teacher, I know how valuable your time is and appreciate your willingness to share it. What is your favorite course? Tell me something about one of your classes that has excited you recently.
3. How would you describe Artificial Intelligence, or AI, to someone who has never encountered the term? (RQ1)
4. How would you describe Machine Learning, or ML, to someone who has never encountered the term? (RQ1)
5. If you do, in what ways do you implement AI and ML topics in your courses?
<If answer is no> In what ways could you implement AI and ML topics into your courses? (RQ2 / RQ3)
 - a. What kinds of AI and ML topics do you currently (or could you) teach in your CS courses? (RQ2)
 - b. What key takeaways do you have (or plan to have) for your students when covering the AI and ML topics? (RQ2)
 - c. What pedagogical, or teaching, strategies related to AI and ML topics do you currently use (or could you use)? (RQ3)
 - d. What challenges related to AI and ML topics do you currently (or could you) face in your CS courses? (RQ2 / RQ3)
6. What are your thoughts on AI and ML being taught in K-12 schools? (RQ2)
 - a. What do you feel is the best approach to K-12 AI and ML education (for example: AI and ML literacy for all students, deeper dive into technical concepts in CS courses, or another approach)? (RQ2)

- b. What do you think should be taught about AI and ML? And at what age? (RQ2)
 - c. What students do you think should be learning about AI/ML? (RQ2)
 - d. What are some issues that come up for you when thinking about teaching AI and ML to K-12 students? (RQ2 / RQ3)
- 7. I'm going to share an example unplugged activity with you and ask some follow up questions. (I demonstrate the lesson by acting as the teacher while the participant acts as a student. I then ask the following questions:)
 - a. What are your thoughts on this activity? (RQ3)
 - b. How do you feel this activity would be received by your students? (RQ3)
 - c. Do you feel unplugged activities (that do not require students to use computers), such as this one, would influence your students' understanding of AI and ML? If so, how? (RQ3)
 - d. What key considerations in the design and delivery of unplugged AI and ML activities would benefit your students? (RQ3)
- 8. There are a lot of topics CS teachers are asked to teach in addition to the concepts covered in the standards (for example: cybersecurity, data science, ethics, among other things). Where would you prioritize AI and ML among these additional topics and why? (RQ2)
- 9. Do you have any other thoughts that we have not discussed that you would like to share?
- 10. *Read closing script* - That concludes our conversation. Thank you so much for sharing your perspective. If you have any other thoughts you would like to share or any questions or concerns arise, please contact me. From here, you will be emailed an Amazon gift card as a thank you for your time and insights. I want to be sure that I am correctly interpreting teachers' responses, so as the study progresses, I will reach back out to you with preliminary findings that I would welcome your feedback on. Thank you for the work you are doing.
- 11. Procedures
 - a. The researcher will ask probing follow-up questions such as, "Can you tell me more about ...?" or "Will you please clarify what you mean by ...?" when necessary to seek more detail.
 - b. The interview will advance to the next question only when sufficient detail has been provided.

- c. If a participant answers a future question in the course of responding to the current question, the researcher will provide a synopsis of their response when they progress to the future question and ask if they would like to elaborate on the synopsis.

APPENDIX E: INTERVIEW UNPLUGGED LESSON PRESENTATION

Figure 1: Slide 1




Figure 2: Slide 2

Lesson Summary

Students develop classification rules from a series of sample elements (training data). The resulting model is used to classify new examples (test data) and the accuracy of the predictions is determined.

Figure 3: Slide 3

Context

We are animal caretakers on a farm. The cows look cute, but some of them bite. We know which ones bite humans; however, new cows will be joining the herd soon and we need to figure out which new cows will bite and which won't - preferably without getting too close to their teeth!





Figure 4: Slide 4

Develop Your Rules



Using what we know about our current cows (training data), think of rules (classification model) to sort cows into biting and non-biting categories.

Your rules could be expressed in text, a decision tree like the one to the left, or in another form you prefer. Be sure that someone else will apply your rules without explanation.

Figure 5: Slide 5




Figure 6: Slide 6




Figure 7: Slide 7

Class Activity

- Split students into groups
- Each group agrees on rules to classify cows as “biting” or “non-biting” based on training data
- Groups exchange rules and apply them to classify each new cow from testing data
- Teacher reveals whether each test cow is “biting” or “non-biting”, groups track their results, calculate the accuracy of their model, and compare results
- (Optionally) Class determines the best model and applies it to classify another group of more challenging test cows

Class Discussion

- Have students explain the rules they developed and point out how each set of rules differed
- Point out that (likely) none of classification models were 100% accurate
- Have students describe their own learning process and then compare it to that of a computer
- Ask students to share examples of classification models they have encountered (examples might include: will this web shopper buy something, is this email message spam, has this picture been Photoshopped, etc.)

Figure 8: Slide 8

Key Takeaways

- In supervised learning, the AI system observes patterns in training data and learns which patterns are typical for each category
- Certain combinations of characteristics indicate an appropriate category
- These rules are used to classify new elements into categories
- Test data is used to determine the accuracy of the classification model
- AI uses the classification model that best fits the given data
- No classification model is perfect
 - Impact of overweighting or neglecting certain characteristics
 - Large training sets can help but can lead to overfitting and inflexibility to new data

APPENDIX F: CODEBOOK

Table 7

Codebook

Category	Subcategories	Sample Codes	Explanation	Example Passage
AI and ML Understanding / Perspectives (21 / 130)	Perceptions of AI and ML and Their Abilities	<ul style="list-style-type: none"> ● dangerous ● magic ● sci-fi ● useful automation ● human-like abilities ● prediction ● surmises user's intentions 	Subcategory contains capabilities that participants explained AI technologies possess and attitudes that participants, their students, or society has about AI technology	<p>“When I think of artificial intelligence, I think of like, very sci-fi stuff”</p> <p>“It [AI] looks like a magic trick. It looks just like like a black box, the black mystery box that does some strange thing, unless you actually know what's really going on.”</p> <p>“But at our level, it's just like, you know, it's [AI] not a movie. So it's not going to blow up people and shoot them or kill them.”</p> <p>“I think a lot of people have like, an image that AI is like the Terminator or something. And it's a very different thing from like, what real life actually is.”</p> <p>“AI is a way to have a computer automate tasks”</p> <p>“I think of it [AI] as trying to get computers to behave the way humans do”</p>

Table 7 (cont'd)

				<p>“Algorithms and machine models that are that could that can think intelligently, like humans.”</p> <p>“Piece of machinery that understands what you want, just by you giving it some information”</p> <p>“Where it might respond to, you know, the input, but then based upon that input, if the customer is not satisfied, they might ask the same question in a different way. So then they're formulating a different answer. And the computer's essentially figuring out, well, maybe I should have done it this way to begin with.”</p> <p>“Why does artificial intelligence need to be limited to what human intelligence is capable of?”</p> <p>“I don't know, I don't have that capacity to actually explain in a way that someone, like just a regular old friend of mine could understand.”</p> <p>“Yeah, I don't know. I don't have like a set definition I go to anything yet.”</p>
--	--	--	--	--

Table 7 (cont'd)

				<p>“I think, in my mind, and you know, I'm still kind of working through some of these ideas.”</p>
	<p>How AI and ML Work</p>	<ul style="list-style-type: none"> ● algorithm ● improves itself ● inputs and outputs ● learn from mistakes ● mathematical technique ● pattern ● thinking for itself ● human does not provide all knowledge ● supervised ● unsupervised ● tuning optimizing 	<p>Subcategory contains participants' thoughts on what makes AI and ML technologies function including how they are trained</p>	<p>“It's more like an algorithm or a different classification of algorithms that can solve problems by simply just doing a ton of data analysis really, more than anything else.”</p> <p>“Machine learning is the is a data analysis process that uses mathematical operations to identify patterns.”</p> <p>“I mean, it's really, it's fancy math, that can take input and create output.”</p> <p>“It's the process of giving the computer training data so that it can then make choices on its own based on that information.”</p> <p>“It [AI] has like a large database, and it starts trying to predict your outcome based on the patterns that is seen in this database.”</p> <p>“AI is something that can be iterative, and build upon itself with no additional human input.”</p>

Table 7 (cont'd)

	<p>AI Types</p>	<ul style="list-style-type: none"> ● ML classifier ● ML is subset of AI ● neural network 	<p>Subcategory contains participant explanations that identify ML classifiers or neural networks as types of AI or that ML is a component of AI</p>	<p>“And that's, you know, how the machine learning process works is that it tries to find patterns. And that's just one form of AI.”</p> <p>“I feel like I use the terms [AI and ML] interchangeably. I don't know if that's accurate, though. I guess, because my understanding of artificial intelligence is that I guess currently, I think it is limited by machine learning.”</p> <p>“So the obvious answer is like, let's say you have images and you're trying to classify them as one type or another type.”</p> <p>“So machine learning is sort of the system that enables image recognition.”</p> <p>“For example, something like natural language processing, which would be something like a chatbot where you could chat back and forth with a program online that could answer your questions.”</p> <p>“I say, “Here's the Google Teachable Machine, train it. This is what a neural network looks like. And it's these orange and blue dots. And nobody knows what</p>
--	-----------------	---	---	---

Table 7 (cont'd)

				<p>the heck is going on. And we're adding neurons and we don't know what any of this means." And I tell them, that's okay."</p>
<p>Importance of Students Learning AI (23 / 181)</p>	<p>AI Knowledge is Necessary For Today And Future</p>	<ul style="list-style-type: none"> ● AI here to stay ● AI is valuable to learn ● AI not all-knowing ● jobs ● application of coding skills to AI 	<p>Subcategory contains participants' arguments that students deserve to learn AI since knowledge of them is necessary for student success</p>	<p>"The reality is, is that it's [AI] here and it's here to stay."</p> <p>"I feel like it's absolutely important to teach the upcoming generation of kids who are probably going to be around our [teachers'] age when AI is ubiquitous and everyone's using it."</p> <p>"There'll be a lot of career opportunities, I think, too, if you have some level of literacy there."</p> <p>"I have kids that are already interested in like, how that works on the development side, like how they could create their own AI tool."</p> <p>"The other part to it that I think would be so beneficial is that it is an application of the skills that they're learning in the coding class. So they're doing these coding classes, they're learning coding stuff, and they're working on stuff and you kind of see the benefits, but take that into an AI thing, it's like, 'Okay, now you</p>

Table 7 (cont'd)

				<p>can take those skills, and actually develop something that's being used.' I think that would be a great thing for applying the skills that they have been developing.”</p> <p>“The job market that these kids are going to be in is going to, is going to be one that leverages AI and AI tools for, you know, various different things. So that you can't see we're doing our students a disservice if we're putting them out there into the world of you know, being an adult and trying to find a good job for yourself and we haven't taught them anything about it.”</p>
	<p>Citizenship</p>	<ul style="list-style-type: none"> ● AI literacy ● appropriate use ● digital literacy ● black box ● demystify ● empowerment ● students influence AI ● copyright infringement ● data collection 	<p>Subcategory contains participant suggestions for students to learn about AI literacy, appropriate use, and be empowered to recognize flaws and work to change AI for the better</p>	<p>“They're gonna have to be somewhat AI literate, when they go out into the world, even if they are just a regular citizen, I think that would be valuable.”</p> <p>“They need to know what it's good for acceptable and responsible use of it [AI]. So I think we would honestly be doing kids a disservice if we didn't have meaningful conversations about that.”</p> <p>“At the level that we understand from working, you know, the level that we</p>

Table 7 (cont'd)

				<p>understand, you know, the way electricity powers our homes, and the way we, you know, understand that there are laws in the country, to me AI is sort of going to be a foundational aspect, like the way we understood, like, we all could probably generally explain how cars are put together on an assembly line, even a lot of the kids could, and no reason they shouldn't be able to do that with AI.”</p> <p>“I think it's important, I think it's, it's [AI] a technology that's out there, it's being used, we're all victims of it, whether we know it or not. So to learn it would only make us better digital citizens.”</p> <p>“I wouldn't feel very comfortable with them just finding it [AI tools] on their own and not having any guidance at all. We have enough problems with kids finding things on the internet and having no guidance.”</p> <p>“I would want students to understand just how much data is collected with every single thing that they do, how that data can be used responsibly or misused.”</p> <p>“What information would be responsible</p>
--	--	--	--	---

Table 7 (cont'd)

				<p>to give to an AI with what do you know what they do with that information? Some of that more like, you know, your internet footprints forever kind of concept, but just in a different way? What does it mean to have your information put in a black box that probably no human touches, but it's still utilized by that proprietor of the AI?"</p> <p>"Getting a better understanding of the devices that they use all the time"</p>
	<p>Awareness Of Potential Harms in Technology Design</p>	<ul style="list-style-type: none"> ● bias ● ethics 	<p>Subcategory contains participants' arguments about the importance of learning about AI's shortcomings and the damage they can cause</p>	<p>"And there's, there's nothing that is without bias. And so we can't blindly trust an algorithm, even though I think a lot of people think you can. But we just, it's unreasonable to trust an algorithm. Because the person, the people that are involved have biases, and usually their unseen biases are unacknowledged"</p> <p>"There can be potential problems with machine learning, depending on like, you know, it's always only as good as its training set."</p> <p>"I think bias is really important. I think being to being able to critically read about things and to and if you were</p>

Table 7 (cont'd)

				<p>involved, to be able to critically think about okay, What are possible biases that are creeping in? And where are we not expanding the pool of information that's being used in the pool of people that are being used.”</p> <p>“Ethics, definitely understanding how AI is currently biased. And that it will be, it will be their responsibility to make sure that it doesn't get worse. And for that, they have to understand why it is biased, how it's biased.”</p> <p>“[AI] Can only be as good as the person who creates it.”</p> <p>“I hope they leave with the understanding that junk in junk out basically, like, you know, it's, it's how well these things are going to perform is totally totally up to how well they're designed and what, what data is being used and the engineer or our computer scientists that's developing it.”</p>
<p>AI Learning Goals (23 / 262)</p>	<p>AI Literacy</p>	<ul style="list-style-type: none"> ● AI tool use ● black box ● demystify ● exposure 	<p>Subcategory contains participants' arguments for</p>	<p>“It [AI] looks like a magic trick. It looks just like, like a black box, the black mystery box that does some strange thing, unless you actually know what's</p>

Table 7 (cont'd)

		<ul style="list-style-type: none"> ● recognition of AI ● understand AI ML in technology 	<p>students learning about AI literacy</p>	<p>really going on.”</p> <p>“Getting a better understanding of the devices that they use all the time”</p> <p>“People need to understand how this [AI] works, and what it does and what it can do for us, and how it can be used.”</p> <p>“Knowing how the solutions work and how the problems work is knowing how to pick the right solution for the right problem.”</p> <p>“By the time a student reaches high school, they should have a good definition of what AI is. And they should know, different tools. That's pretty much the baseline of literacy of AI.”</p> <p>“If they're interacting with the tools, then they can learn about them. And they're not necessarily going to be learning about them at nitty gritty level. But what is it? When has it been frustrating?”</p> <p>“So when they start understanding, this is not magic, and there's something that can be done, or they have a level of influence and can talk about it. That</p>
--	--	---	--	---

Table 7 (cont'd)

				<p>helps empower them to understand the situation and maybe what they can do about it.”</p> <p>“Even if our generation ruins the whole thing, you know, they, they still have it in their hands to, to fix some stuff, and to start now, to try to make a difference. And if they don't understand how it works, if it's if it's magic, then they're not going to be able to do anything about it, it's going to continue being dangerous.”</p>
	<p>AI Ethics</p>	<ul style="list-style-type: none"> ● bias ● empowerment ● students influence AI ● copyright ● data collection ● ethics ● only as good as training data ● representation in design ● responsible design ● unintended consequences 	<p>Subcategory contains participants' arguments for students learning about AI ethics</p>	<p>“It would be advantageous to get instruction to students about how to use it [AI] ethically.”</p> <p>“The whole point of artificial intelligence is writing programs, which so often is kind of not seen by the general public, but it's programs in the computer, that are written to have the computer generate its own information.”</p> <p>“I think that's the important part that they should understand that that the training of the AI is still done by humans”</p> <p>“And just that the most important part of any development is the people who are</p>

Table 7 (cont'd)

				<p>going to be using the development and making sure that their concerns, their, not their concerns, but making sure their humanity is like being factored into the development.”</p> <p>“We accidentally didn't think about this other group that we're biasing against, without really meaning to.’ Like you put together this well-meaning algorithm, but you didn't look at your data to see if it had bias in it.”</p> <p>“What ways can you impact or change the AIs that are out there?”</p> <p>“My hope is always that students will maybe fix something that is broken.”</p>
	<p>AI for Learning</p>	<ul style="list-style-type: none"> ● AI not all-knowing ● appropriate use ● consider plausibility ● intelligence augmentation ● problem solving ● technique-problem fit ● AI raises the level of assignments 	<p>Subcategory contains participants' arguments for students learning about AI to enhance their learning</p>	<p>“This thing [ChatGPT] is actually really good at finding sources ... you can definitely use it to find sources and then go do your own research.”</p> <p>“This [large language models] is something that's helpful for your [students'] research.”</p> <p>“We sort of like to study kind of the limitations of it [AI], and how it's not a, it's</p>

Table 7 (cont'd)

		<p>and thinking</p> <ul style="list-style-type: none"> ● AI tools to meet student needs 		<p>not a cure all.”</p> <p>“I would even encourage them to say, ‘Cool, here are some AI generated solutions for this problem. Which one of these makes the most sense? If one of them does, does it even fit our problem?’”</p> <p>“There's this big looming elephant in the room of, ‘Okay, these tools are being used to cheat.’”</p> <p>“Teach them how to use it [AI tools] appropriately, rather than cheating with it.”</p> <p>“With my students, I have asked them to check into it [AI], you know, because I want them to be aware and to understand I don't want to put blinders on their eyes. I want them to use it responsibly.”</p> <p>“There could be versions of AI that are designed for interaction with those with either limited language skills or with early language skills”</p>
AI within K-12	AI Literacy for	● AI all students	Subcategory	“I think since this [AI] is so central to the

Table 7 (cont'd)

<p>Curriculum (23 / 143)</p>	<p>All</p>	<ul style="list-style-type: none"> ● all grades K12 ● core subject integration ● digital literacy ● exposure ● K8 AI instruction challenge ● age of maturity ● AI needs to be in standards ● CS not required ● cybersecurity ● don't know best age ● elementary conceptual 	<p>contains participant passages promoting AI instruction for all K-12 students</p>	<p>technology space that we're in now, it would be good if they [all students] got just a little bit of experience, if everyone got a little bit of experience in it.”</p> <p>“I think at least exposure is going to be super important, you know, maybe a month or something of exposure.”</p> <p>“I think yeah, having, having the information for, for all students would be the best approach. Yeah, not every student is going to be technically like, inclined to go down the technical path. Although I think I think there's advantages to to it, just integrating pieces of it for for all students”</p> <p>“Yeah, not every student is going to be technically like, inclined to go down the technical path. Although I think, I think there's advantages to, to it, just integrating pieces of it for, for all students, I mean, if you like I have a programming assignment for my geometry students.”</p> <p>“It would be fun to learn how to integrate it into other courses”</p>
----------------------------------	------------	---	---	---

Table 7 (cont'd)

				<p>“I definitely think that it would be an ideal world if it was addressed or in other subjects”</p> <p>“It'd be good in the lower levels, to just have an introductory, this is what machine learning is, this is what AI is, this is generally what's happening when we do that”</p> <p>“Well, as far as the younger kids, I don't think it'd be easy to have them do a lot of hands on because it uses it uses so much, use a lot of programming skills that they may not have yet. So I would want them to mainly have that literacy at a younger ages.”</p> <p>“I think that there's no reason why the intricacies of machine learning would have to be taught in high school. But I think AI literacy, yes”</p> <p>“So I think in the lower levels, if there was more awareness of what programming is computer science, that kind of thing. So that as they get older, they can make better decisions in middle school to for their electives that they take there.”</p>
--	--	--	--	---

Table 7 (cont'd)

				<p>“As of right now, it's non existent, in least in our standards, and <State> is not known to be the fastest adopting technological standards”</p>
	<p>AI in Elective High School CS Courses</p>	<ul style="list-style-type: none"> ● AI opt in ● not every student interested in technical side ● stand-alone AI course ● student interest ● AP structure challenge ● secondary programming 	<p>Subcategory contains participant passages promoting AI instruction in elective K-12 courses since not every student will have an interest in the subject</p>	<p>“Developmentally and interest wise high school would make sense to then jump into some of the more technical pieces [of AI], maybe for all students to have a basic understanding or for those who really are interested in, in the details that could come in a CS course or somewhere.”</p> <p>“I get students now who don't actually opt in, they're just put in to some of my classes, and they have very little interest at all, in any of that kind of stuff ... so it doesn't meet their needs at all.”</p> <p>“My kids were kind of forced into taking a coding course, they didn't like it”</p> <p>“I'm thinking of student of my students and some of them will not be interested in the technical side of it, like, some will, because some are intuitive, though, they'll want to know how it's built, how it's created, how its programmed to learn</p>

Table 7 (cont'd)

				<p>from itself and from other things. But I think mainly as a as a kind of like a general thing. I think it's mainly the exposure and the ethics to it. Because not every student is looking for that technical knowledge.”</p> <p>“I mean, foundationally, I guess it would be nice if like AP, actually could be a little more nimble and actually see, you know, yes, it's important to teach all these things about Java, but like there's other important pieces that are being left out entirely. And maybe they ought to update a little.”</p>
<p>AI Pedagogy (20 / 78)</p>	<p>Student Centered Pedagogy</p>	<ul style="list-style-type: none"> ● encouraging experimentation ● no lecture ● personalized learning ● programming ● project based ● real world ● service learning 	<p>Subcategory contains participant passages detailing student-centered pedagogies they feel are valuable when teaching AI topics</p>	<p>“So sometimes just experimenting and playing with different things, and just being, having fun with the frustration has been helpful.”</p> <p>“So I don't believe in standing and lecturing at kids ever. I feel like that mode of teaching doesn't work very well, especially for something that's a very practical skill that each person kind of has an individual relationship with as they start to learn how to work with it.”</p> <p>“So in a lot of ways, the reason I do the</p>

Table 7 (cont'd)

			<p>flipped classroom and like, try to individualize the lessons closer to each individual person is because each person is different. And so I'm trying to find the issues they have as fast as I can and make sure that we find a way that they can learn as quickly as they can for their particular learning style, and their particular background.”</p> <p>“My belief is that by working through a project, working through ambiguity, I use that word a lot with my students from the beginning, being curious and using ambiguity to learn. And so projects, I think, are the greatest way for them to kind of confront that ambiguity, often with somebody to support them up here, and to have enough time for each person to work through those projects”</p> <p>“I love to give students project-based. I'm really huge into allowing them to be investigators, explorers of the information.”</p> <p>“I want the students to do a service project where they ask the AI what are some societal issues that we can that we can address as high school students”</p>
--	--	--	--

Table 7 (cont'd)

				<p>“Socratic seminar type discussions, and having some of those kind of student driven open ended top conversations around AI and maybe it would be helpful to have like, some case studies or some like live examples that we could run real quick.”</p>
	<p>Teacher Led Pedagogy</p>	<ul style="list-style-type: none"> ● collaborative file editing ● current events ● discussion ● readings ● societal issues ● video 	<p>Subcategory contains participant passages detailing teacher-led pedagogies they feel are valuable when teaching AI topics</p>	<p>“Giving students time and space and examples of how that happens, not just in technology, but kind of just in life. Because I think that can be very, very hard for students to conceptualize. I think examples are probably like, the most useful tool”</p> <p>“Using some of the current news cycle.”</p> <p>“When news comes up I bring it to them as kind of an enrichment thing.”</p> <p>“I think AI ends up a little more in the philosophy world than a lot of people expect just because there are a lot of big implications in it. So I do, probably do a lot more kind of thought pieces. ‘And what is your reaction to this? And what is the real world consequences if you were to follow through on this?’”</p>

Table 7 (cont'd)

				<p>“Examine some case studies about implementations of AI. But then to kind of build some as well, or some build some, some plans, project plans for implementation.”</p> <p>“This curriculum, they've provided a lot of videos, you know, so I would just, thought I'd show them this video, and then I would, you know, have some questions. And then some discussion after I showed them the videos.”</p> <p>“I showed three different videos, one on self driving cars, Waymo, is that Waymo? One on Neuralink's work on brain implants, and one on Chat GPT”</p>
	<p>Student Tools for Instruction</p>	<ul style="list-style-type: none"> ● no/low code tool ● robotics 	<p>Subcategory contains participant passages that list technology tools they have found useful when teaching AI and explain the benefits</p>	<p>“Well, a lot of the no code tools helped me with that, because they get to kind of play with stuff. So we're not just doing all this heavy, heavy talking, we're also doing a lot of playing and keeping things fun.”</p> <p>“That's why I love block programming. But because we can, we can do all the computational thinking [of AI algorithms] with none of the silly typographical errors</p>

Table 7 (cont'd)

				<p>that a lot of my students make [when they write programs using text-based programming languages].”</p> <p>“I prepare Jupyter Notebooks that I share with them. And then they're, they're able to see that there is some code behind, but they don't have to look at the code. And they can just click, and maybe fill in some text boxes”</p> <p>“They could start learning about how the robot could teach itself, you know, could could do some machine learning and it, you know, be an AI in that sense”</p> <p>“Anybody's going to associate AI and machine learning with robotics. Because the more the robot can teach itself, without having to be explicitly programmed in force, make them more useful.”</p>
	Class Communication Norms	<ul style="list-style-type: none"> ● class norms ● communication 	Subcategory contains participant passages that point out the importance of class norms to	<p>“We do talk in the beginning about conversations can become uncomfortable, and that this needs to be a safe place. And, and what kind of the class itself comes up with? What's a code of conduct? How do we handle ourselves when we talk about things that</p>

Table 7 (cont'd)

			<p>develop a safe classroom discussion environment and the benefits of students building communication skills</p>	<p>are difficult if we become uncomfortable?”</p> <p>“The point is not to make people uncomfortable, but is to talk about uncomfortable situations so that they are prepared when they go out in the world to have those conversations with people who are maybe not in such a safe place.”</p> <p>“I think it's very helpful for them to kind of work together in some smaller groups like that so that they're actually talking to one another.”</p> <p>“They also have jobs that they have to do in the groups of four. So forcing them to talk to each other, and communicate.”</p> <p>“I force my nerdy computer science geeky kids into having to talk to their neighbor, which they would prefer to never have to do. And that's why they signed up for computer science in the first place, thank you very much. And so it forces a little bit of interaction that they're naturally disinclined to do”</p>
<p>Unplugged AI Activities</p>	<p>Benefits Of Unplugged</p>	<ul style="list-style-type: none"> challenge assumptions 	<p>Subcategory contains</p>	<p>“Having something early on that makes them stop and do a little more thinking</p>

Table 7 (cont'd)

<p>(19 / 73)</p>	<p>Activities</p>	<ul style="list-style-type: none"> ● helps AI understanding ● many right approaches ● no programming expertise necessary ● think in new ways ● true collaboration 	<p>participant passages explaining the benefits of unplugged AI activities in the K-12 classroom</p>	<p>than just going with their instinct, in that in an activity is also the thing”</p> <p>“It's understanding the concepts first, and there's no pressure on knowing the technology of it.”</p> <p>“I feel like it's appropriate for any level of students, I feel like you could do that in, you know, you could do it in a sixth grade class, you could do it with your seniors at the end of the year, you could do it with anybody.”</p> <p>“I really think they're better than just going straight to the computer for actually understanding the process of how machine learning works, you can't you can't get the process just by watching the tool work more”</p> <p>Students are “trained and taught that there is a right answer to things that there just is a correct way to do it. And I think anybody who's ever taken any more than just an introductory class in computer science knows that there's actually a million different ways to get to the right answer, and honestly, usually, it's through a lot of wrong ways in the first</p>
------------------	-------------------	--	--	---

Table 7 (cont'd)

				<p>place.”</p> <p>“It helps because it's not code. So they're not they're not trying to figure out the code and see how that they're not bogged down with the code. They're, they're really just experiencing what, what is going to happen when they do AI and how it works without any kind of code.”</p> <p>“Coding sort of like, sort of similar to math, some students just kind of tense up at the thought of it. And there's this kind of preconceived notion that the student will have that they're not going to be good at it.”</p> <p>“But in terms of being able to have those discussions without having to write the stuff behind it, it's really important, I think.”</p> <p>“Some of these unplugged activities really helped them to think more. And think at a higher level than when they are just on the computer doing stuff”</p>
	Keys To Unplugged	<ul style="list-style-type: none"> ● competition ● connection to 	Subcategory contains	“Really gamifying things really gets them intrigued and engaged.”

Table 7 (cont'd)

	<p>Success</p>	<p>students' lives</p> <ul style="list-style-type: none"> ● focus on topic ● hands on ● limited vocabulary ● simplicity ● student creativity 	<p>participant passages explaining the important elements for successful unplugged activities in the K-12 classroom</p>	<p>“If you can figure out how to take that thing and tie it to something that's real in their life. I think it could be very powerful.”</p> <p>“it's very hands-on something that the students can can like, grab with their hands, and they the, the build it themselves, they test it out themselves, like physically, I think anything that has to do with that, in a computer science class is always like very refreshing for students, because they're not stuck in front of a screen.”</p> <p>“Doing it as an Unplugged Activity is so much better because they're doing it, as opposed to just you know, hearing about it. Okay, so this is what does, it doesn't really, you don't really reflect on what that really means to come up with it. But if you go through those steps of that activity, you're actually having to come up with that model, and then start to think about, Okay, what did I miss in that model?”</p> <p>“If you kind of hide the, the content vocabulary before doing the activity, and</p>
--	----------------	---	---	---

Table 7 (cont'd)

				<p>then bring it back in at the end.”</p> <p>“It has to meet them where they are, it's it has to be something that is very simple to understand to begin.”</p> <p>“They like kind of that ownership and creativity.”</p>
	<p>Challenges Of Unplugged Activities</p>	<ul style="list-style-type: none"> ● plugged sometimes better than unplugged ● logistics challenge ● activity must fit learning environment ● preinstruction 	<p>Subcategory contains participant passages explaining the challenges of unplugged activities in the K-12 classroom</p>	<p>“I feel like setup has to be easy. So in other words, like, as a teacher, it's gotta be very straightforward, you know, I have a 40-minute class period. So you know, so that's not a ton of time.”</p> <p>“My biggest hurdle in using unplugged group activities is dealing with the space concerns.”</p> <p>“I do think I haven't always loved unplugged stuff. And maybe that's could be my bias could be, because it's like a pain in the butt to get kids to do them sometimes. And it could be because I just had to print something out. It's like, Oh, crap. Now to get the printer, you know, all those logistical things that kind of like, take up that extra bit of time.”</p> <p>“It definitely requires a lot of</p>

Table 7 (cont'd)

				<p>preinstruction, right? So they understand what the goal is, like. You give them a bunch of materials and then without explaining exactly what the goal of all these materials are, then it would be difficult, it can be difficult. So really knowing what their goal is ahead of time and kind of also not trying to like a lot of times the unplugged activities seem to be kind of exploratory or investigative.”</p> <p>“Lack a direct application to the topic at hand. Sometimes I feel like the unplugged lessons, they’re a little like topic adjacent.”</p>
<p>Challenges Of Teaching AI (21 / 163)</p>	<p>Teacher Capacity</p>	<ul style="list-style-type: none"> • don't know what should be taught • teacher content knowledge 	<p>Subcategory contains participant passages detailing the challenges around teacher capacity when teaching K-12 AI</p>	<p>“I just don't know what it is that, what it looks like, or what the lessons should be.”</p> <p>“I think for me, it'd be more on my end, where I just don't know enough about it, to, you know, put it in, to actually implement it into the curriculum”</p> <p>“I think the main issue is time”</p> <p>“The biggest challenge is just the amount of time and I'm spread so thin”</p>

Table 7 (cont'd)

			<p>“I’m trying to think of so one of the caveats I must say though, is that like teachers are already at like a zero sum game. So what is this lesson replacing that they’re already doing? Because you have to do that. There’s no more time.”</p> <p>“It’s just finding the time. So it’s not only not having enough time to teach all the concepts in your course, it’s for you for your own professional learning sake, to prepare.”</p> <p>“I think we have to start with fundamentals, because I’m not familiar so I want to know about those things, too.”</p> <p>“I don’t have all that knowledge myself right now. So saying how I would teach that I don’t really know”</p> <p>“I want to make sure I know enough of it to be able to teach it appropriately, and make sure I’m not misleading my students. So I’m a little bit worried about just kind of getting a good mastery of the content myself.”</p> <p>“Not having time on my own and to to</p>
--	--	--	---

Table 7 (cont'd)

				<p>learn everything that I need to learn.”</p> <p>“Where are they gonna get the people to teach these things? I'm serious. That's like, it's like putting the horse before the carriage. You know, we are really, really short.”</p>
	Student Ability Limits	<ul style="list-style-type: none"> ● collaboration challenge ● limited ability to carry out complex project ● limited math background ● limited programming experience 	<p>Subcategory contains participant passages detailing the challenges with students' pre-existing ability levels when teaching K-12 AI</p>	<p>“A lot of machine learning is very procedural work. And, and the, the amount of focus that's needed for that, I think there's just a few students who have the interest and the focus.”</p> <p>“Trouble there is that their mathematics level is not high enough.”</p> <p>“The math part is challenging, because if we want them to understand how a classifier works, sometimes we have to go deep into math that they don't they have not learned yet.”</p> <p>“Yeah, I probably wouldn't have everybody in the school take an AI computer science course, you know, where they're actually programming. Because I think that will require a certain skill level.”</p>

Table 7 (cont'd)

				<p>“Yeah, I probably wouldn't have everybody in the school take an AI computer science course, you know, where they're actually programming. Because I think that will require a certain skill level.”</p> <p>“I always survey them [incoming students], ‘What have you ever done before?’ Rarely have they ever done any programming. And so it's all new.”</p>
	<p>Traditional Structures and Processes</p>	<ul style="list-style-type: none"> ● administrator challenge ● AP structure ● engrained pedagogy ● CS teacher limited influence on what's taught ● parent resistance ● stereotype ● time 	<p>Subcategory contains participant passages detailing the challenges around overcoming established culture and tradition in order to teach K-12 AI</p>	<p>“Administrators don't really know what you're doing ... I mean, most administrators these days don't even know what computer science is. ... But unfortunately, sometimes they might not even know enough about it to value it.”</p> <p>“AP Computer Science A would be super hard, just because we're pretty tight on everything that we have to learn. So to veer off and do that might be a little bit that would be the time part”</p> <p>“Because of the power that AP has, you know, that they're kind of controlling what students are being taught all over the country. And that's apparently important for I don't know, if students feel it's</p>

Table 7 (cont'd)

				<p>important, and colleges feel it's important, this whole AP thing. So there's a lot of power there. And if that doesn't change, it's going to be hard to get, like, AI to be considered an important topic covered in school, unless it's some other non AP computer science course.”</p> <p>“There's no way I can fit it into AP CSA”</p> <p>“There are a lot of teachers that have a way of teaching, and honestly, they will work 10, 15, 20 years, find a way that works and in their mind and in their lives it is how you teach and there's not a different way.”</p> <p>“The community in which they're teaching is saying, no, no, we need to ban this from our schools period. For whatever reason, and then students are caught in this like area of not either getting the information they need to understand how it works, or getting the boundaries on like, these are the ethical ways in which you can use it”</p> <p>“But still there is just there's some stereotypes you fight against. And so all of those are the things that you're really</p>
--	--	--	--	--

Table 7 (cont'd)

				<p>working to try to combat as you're trying to get more people into interested in it into it.”</p> <p>“The my goal would be to use it in an assistive manner to be to do what computers already do for us to take what we can already create and make it something that's a lot more efficient.”</p> <p>“Okay, well, you have a calculator. So now, now, you don't need to teach people to memorize the multiplication table anymore. You know, I said that people don't do long division by hand anymore. But you are now free to do other things and more sophisticated things.”</p>
	<p>Limits of AI Curriculum and Learning Tools</p>	<ul style="list-style-type: none"> ● adapt curriculum ● AI in adopted curriculum ● curriculum not at appropriate level ● lack of curricula challenge ● third-party curriculum ● value of curricula 	<p>Subcategory contains participant passages detailing the challenges around choosing, creating, and adapting curriculum to teach K-12 AI</p>	<p>“I like unplugged activities. And like I said, it's just a matter of having them available to use.”</p> <p>“I just haven't found a lot of lessons on artificial intelligence.”</p> <p>“That's the hardest part, you know, for these schools that are starting up is they got is finding good resources.”</p>

Table 7 (cont'd)

		<ul style="list-style-type: none"> ● access to AI tools ● lack of appropriate tools 		<p>“Availability of tools and sort of starter curriculum, I guess, would be the summary of the challenge that I'm seeing.”</p> <p>“Just some one-stop shopping for teachers like, Here are the options for you. Maybe something that's vetted by, you know, people, other teachers some way to, like, communicate this in a streamlined way.”</p> <p>“But you know, having those kinds of things because teachers just don't have the time to put this together. They're not experts in this field. And, you know, having those kinds of things available would be really helpful in helping it get adopted in schools.”</p> <p>“ I did put together a lot of curriculum, so, I'm not opposed to that. This field in particular seems like it's changing so quickly. And, you know, I'm not saying that I have a deep understanding of it. So that's why it would be somewhat reliant on, you know, getting some of the experts to put it together, not just experts in AI, but experts in AI and pedagogy,</p>
--	--	---	--	---

Table 7 (cont'd)

				<p>you know, putting it collaboration, that would be good.”</p> <p>“I’m always looking for activities that is more at their level. So that to expose them to it without having crazy assignments, that would be, you know, university level college level assignments.”</p> <p>“The actual work was, it was like, right in that that terrible middle spot where it was way too hard for some kids, it’s way too easy for others.”</p> <p>“Our school district actually has completely restricted access to the website on their networks ... the students actually currently have no way to use them outside of their cell phones and cell phone service.”</p> <p>“We can’t access anything unless it’s vendor approved or COPPA approved, you know, and so that becomes a challenge.”</p>
--	--	--	--	--