International Journal of Artificial Intelligence in Education (2023) 33:384–398
https://doi.org/10.1007/s40593-022-00327-w

ARTICLE

Designing Game-Based Learning for High School Artificial Intelligence Education

Maxyn Leitner[1] · Eric Greenwald[2] · Ning Wang[1] · Ryan Montgomery[2] · Chirag Merchant[1]

Accepted: 15 December 2022 / Published online: 10 April 2023

Abstract

Artificial Intelligence (AI) permeates every aspect of our daily lives and is no longer a subject reserved for a select few in higher education but is essential knowledge that our youth need for the future. Much is unknown about the level of AI knowledge that is age and developmentally appropriate for high school, let alone about how to teach AI to even younger learners. In this theoretical paper, we discuss the design of a game-based learning environment for high school AI education, drawing upon insights gained from a prior cognitive interview study at a STEM focused private high school. We argue that game-based learning is an excellent fit for AI education due to the commonality of problem solving in both game playing and AI.

Keywords Artificial intelligence · K-12 education · Game-based learning

✉ Eric Greenwald
eric.greenwald@berkeley.edu

Maxyn Leitner
maxynrl@usc.edu

Ning Wang
nwang@ict.usc.edu

Ryan Montgomery
rmontgom@berkeley.edu

Chirag Merchant
merchant@ict.usc.edu

1 University of Southern California Institute for Creative Technologies, Los Angeles, CA, USA

2 University of California, Berkeley, Berkeley, CA, USA

Introduction

How can we best prepare K-12 students as citizens of a world dependent on Artificial Intelligence (AI)? AI is increasingly inescapable in our world, permeating an ever-growing list of domains. Medical diagnosis, driving cars, parsing human speech, and even suggesting what we should watch in our free time are all tasks that AI is becoming ever more competent at. Rapid advances in the design and implementation of AI systems to accomplish these kinds of automation have led to the ever-expanding role for AI in society (Makridakis, 2017; Nadikattu, 2016). AI, it seems, is all around us.

Accordingly, AI is redefining the future of work within the human-machine alliance (Guszcza et al., 2017). Thus, proficiency in the language of AI is key to a workforce that will continue to innovate and support the AI-powered technology infrastructure. All of today's students will go on to live a life heavily influenced by AI, and many will work in fields that involve or are influenced by AI. It is no longer sufficient to wait until students are in college to introduce AI concepts. Rather, they must begin to work with AI algorithmic problem solving and computational methods and tools in K-12.

At the same time, although AI's impact on society is increasingly pervasive, and innovative educational opportunities are being rapidly developed, there has been previously little research into how students, especially pre-college aged students, construct an understanding of and gain practice with core ideas in the field. As a result, there is little possibility of grounding the design of learning experiences in evidence-based accounts of how youth learn AI concepts, how understanding progresses across concepts, or what concepts are most appropriate for what age levels.

Game-based learning environments are increasingly being explored for their positive effects on both cognitive processes and engagement (Gee, 2003; Wouters et al., 2013). Such learning environments are particularly well suited to teaching students how to problem solve, as gaming inherently elicits constant experimentation and reflection that can be used to ingrain real-world problem-solving skills (Salen, 2007). This kind of approach to education is perfectly suited to teaching AI, as AI is problem solving at its core. AI learning is learning how to approach tasks in a systematic way, in service of developing algorithms that can be run by a computer to solve problems. To this point, game-based learning environments are already being developed to teach AI concepts to upper elementary school students (Lee et al., 2021). For these reasons, game-based learning is an excellent technology with great potential for K-12 AI pedagogy.

We have developed a game-based learning environment aimed to help high school students build AI competency, apply math knowledge, and develop computational thinking (CT) skills. Though this design has not yet been evaluated in any studies with students, our design decisions are informed by common needs of high school STEM students illuminated in a previous cognitive interview study. We proceed to identify important takeaways for the variety of AI students present in high school learning environments. Finally, we detail the design of the game-based learning environment and discuss how it aims to address the aforementioned needs and ensure the desired takeaways.

Related Work

Although AI has been the cornerstone of the computer science curriculum in higher education for decades, discussions on how to approach AI education for the K-12 population have only just begun in the US (Touretzky et al., 2019), Europe, and much of the rest of the world. In contrast, China (Xiong et al., 2018; Chen & Tang, 2018) has already developed a series of seven AI textbooks for elementary, middle, and high schools, and Sweden has also developed AI courses to educate its citizens, including school-age youth, about AI (Heintz et al., 2015). Meanwhile in Europe, the European Driving License for Robots and Intelligent Systems is under development and will soon be implemented as a standardized, internationally accepted system for training and certifying educators and high school students in AI and Robotics (Kandlhofer et al., 2019).

Discussions on how to integrate AI into the existing K-12 curriculum (e.g., computer science education) are heating up in the US (Gardner-McCune et al., 2019). To develop guidelines of what K-12 students should learn about AI, the AI4K12 Initiative has proposed the Five Big Ideas of AI, including Perception, Representation and Reasoning, Learning, Natural Interaction, and Social Impact (Touretzky et al., 2019). Most recently, ReadyAI, a group that organizes camps to help students learn about AI, has developed a curriculum to teach AI courses to K-12 students online at ReadyAI.org. Researchers at MIT have also developed a website to share a variety of online activities for K-12 students to learn about AI, with a focus on how to design and use it responsibly.[1] This includes a curriculum for teaching ethics of AI to middle school students.[2] MIT research in this area has also led to the development of PopBot, a learning experience for children ages 4–6 teaching AI and AI ethics (Ali et al., 2019). Meanwhile, work at the University of Colorado is developing modules integrating AI ethics education into robotics learning (Yeh et al., 2019).

On the technology front, there has been effort, particularly from industry, to build demonstrations and tools to help the public learn about AI, particularly machine learning (for review, see Gardner-McCune et al., 2019). Additionally, Carnegie Learning has developed a prototype to help middle school students learn AI by designing AI to play tic-tac-toe (Ritter et al., 2019), and the Frankie Project aims to utilize a friendly looking robot with shape recognition capabilities to prompt students to discover how neural nets can be trained to perform such tasks (Pimentel et al., 2018). LearningML, a platform where students can experiment with ML models and develop applications that use them, has also seen success in introducing machine learning to students aged 10–16 (Rodríguez-García et al., 2021). However, simply providing students with AI tools is not sufficient AI pedagogy. K-12 AI education requires teaching students how to use such tools.

Game-based learning environments have shown promise in helping students learn the kinds of problem-solving skills and analytical thinking vital to understanding AI and show great potential for this much needed pedagogy. Unlike non-interactive

[1] aieducation.mit.edu/.

[2] www.media.mit.edu/projects/ai-ethics-for-middle-school/overview/.

forms of media, problem solving is at the heart of games. AI is fundamentally about problem solving, making game-based learning environments a phenomenal fit for teaching the subject. PRIMARYAI, a collaborative game-based learning environment integrating AI knowledge with life-science topics, is one such project aimed at teaching AI to upper elementary students in recognition of this need for all students to have opportunities to learn more about AI (Lee et al., 2021). By making AI algorithms the mechanism for solving in-game puzzles, students can learn how and when to apply them to real-world "puzzles." In addition, narrative design in game-based learning environments can encourage students to "play a character," stepping into a different role such as a researcher or explorer that prompts them to learn and apply the concepts being taught (Gee, 2003). Crystal Island, a game-based learning environment designed to teach microbiology to middle school students, has seen success by effectively utilizing narrative design to create a positive relationship between learning and engagement with the game, suggesting that two factors are important in creating this positive relationship (Rowe et al., 2011). First, the narrative, educational content, and gameplay should all be tightly integrated. Second, the narrative should be complex and interesting enough to motivate the student to proceed through the game, but not be too complicated and risk distracting the student from the learning at hand. These principles will inform the design of our game.

The work presented here aims to uncover how to design a game-based learning environment to meet the challenges of teaching AI to high school students, with hopes of building a foundation for future designs aimed at younger K-12 grade bands. This work builds upon explorations into how high school students approach AI concepts, what obstacles they face, and how to guide them through those obstacles (Greenwald et al., 2021). This work also draws upon previous investigations into linking AI to high school math curriculum to identify AI concepts suitable for high school students (Wang & Johnson, 2019), as well as work investigating the learning of computational thinking (Rich et al., 2019) and foundational research into comprehension of mathematical representations (e.g., Curcio 1987; Friel et al., 2001) and statistics (e.g., Batanero et al., 1994).

Methodology

Our methodology first builds off previous work connecting AI concepts to the underlying math knowledge required, suggesting a selection of topics based on whether the relevant math knowledge is covered by the high school curriculum (Wang & Johnson, 2019). Primarily, AI concepts in the field of Knowledge Representation and Planning were excluded due to the prerequisite of knowledge in logic (logic is not commonly part of the K-12 math curriculum in the United States). In the end, five AI concepts were suggested: Search, Bayesian Networks, Decision Trees, Clustering, and Linear Regression. Search can connect with and build upon student knowledge of exponential and linear functions in the exploration of how a search tree grows as depth increases. Bayesian Networks draw upon students' study of Bayesian probability, and Decision Trees rely upon students' knowledge of logarithms. Clustering has a similar connection to high school geometry, and Linear

Regression itself is a subject in high school mathematics classes. Using the classic example problems for the selected AI concepts, we designed AI problems set in scenarios that are familiar to high school students, along with solutions and scaffolding plans.

In previous work, we conducted a cognitive interview study with high school students to better understand the range of background knowledge and experiences students bring to AI learning, the concepts that are most readily accessible or challenging, and the strategies to help students learn AI concepts through leveraging existing knowledge and experiences (Greenwald et al., 2021). These students all attended a private high school in the Los Angeles metro area with a strong focus on Science, Technology, Engineering, and Mathematics (aka STEM). AI concepts to teach were determined based on the AI curriculum from the most popular AI textbook for higher education (Russell & Nerving, 2016), categorizing the AI topics into four main fields (Search, Knowledge Representation and Planning, Probabilistic Reasoning, and Machine Learning). Problems were then designed for students in Search, Probabilistic Reasoning, and Machine Learning.

Each problem consisted of a series of questions that helped students develop the solution step by step. These AI problems were used in the aforementioned cognitive interview study and served as a foundation for the scenario design for the game-based learning environment. During the study, the students were asked to solve AI problems with paper and pencil and were encouraged to think-aloud while they worked on the problems.

Common Needs

The design of our game-based learning environment seeks to address three important needs found across high school STEM students during our previous cognitive interview study: explicit scaffolding to establish abstract representations characteristic in AI problems, explicit scaffolding to parse problems in terms AI systems can operate on, and support to leverage and apply mathematical concepts that underlie AI problems (Greenwald et al., 2021). These needs were considered within the specific context of Classic Search as a topic.

For establishing abstract representations, it was found that students needed explicit scaffolding in understanding how to interpret and construct a search tree based on a search scenario: it was a novel and non-intuitive way to represent the problem space of search for most students. For both pre-constructed search trees and trees constructed by the student with instructor guidance, students struggled to understand how puzzle-solving tasks could be conceptualized as a search problem, how search problems could be represented as trees, and how the computational features of tree representations (e.g., the breadth and depth of the tree) were connected to the complexity of the problem. Each of these difficulties echo recent findings (see, e.g., Basu et al., 2016) related to youth challenges with computational thinking more broadly. Relatedly, understanding how to mathematically estimate time and space complexity proved to be difficult for students: getting traction on making and interpreting these estimates depends on a working understanding of how computers process information and manage computational resources (e.g., what information needs to be stored for different search algorithms and how those needs impact time required to execute a command). These findings evoke long-identified challenges (e.g., Du Boulay, 1986) for CS students in understanding the abstract features of the "notional machine" a program is controlling, and how that impacts the real world. Finally, representing abstract problems such as puzzles in a way that enabled students to apply search techniques to them required significant support. This finding is consistent with earlier cited research in CS education and resonates with well-established findings from mathematics education (Friel et al., 2001) about student difficulty making meaning of mathematical representations. Each of these challenges represent skills that we aim to develop with our game-based learning environment.

Applications of AI Knowledge

In addition to these common needs found across students in the study, previous work discussed how students with different relationships to AI will have different desired takeaways. This comes down to roughly three different ways of interacting with AI as discussed within the K-12 AI education research community (Gardner-McCune et al., 2019): those who will be AI end-users in their everyday lives and careers, those who will need to implement AI algorithms, and those who will research and advance AI algorithms. We refer to these categories as AI Consumers, AI Operators, and AI Developers respectively, though others may use the terms Users, Implementers, and Researchers.

AI Consumers refers to students who will consume AI in their everyday lives, i.e., all students. For these students, it is important that they learn what is possible with AI, and what the impact of AI will be on them. For example, understanding how AI might be used to help diagnose disease or design targeted advertising campaigns will allow these students to better grasp what advances in AI mean for them personally.

AI Operators refers to students who will have to make use of AI in their careers and will have to make informed decisions about how to utilize it effectively. These students will need an understanding about the variety of AI algorithms, the technologies that they use in their work, and what their strengths and weaknesses are. One example might be a student who goes on to work in the biomedical field, and thus makes use of AI technologies in order to analyze data. Understanding the algorithms underlying the technologies that they are making use of will better enable them to select the technology best suited to the situation at hand.

AI Developers are the ones who will go on to develop new AI technologies. For these students, a much deeper understanding of AI algorithms is necessary. These students must know the ins and outs of each algorithm, to not only know what its strengths and weaknesses are but how the design of the algorithm resulted in those strengths and weaknesses. They must know how to estimate space and time complexity for these algorithms and understand what implications these quantities hold for the effectiveness of a particular algorithm in a particular scenario.

ARIN-561

The game-based learning environment we have designed, ARIN-561, aims to help students learn AI concepts, apply math knowledge, and build computational thinking (CT) skills. In the game, students play the role of a scientist who sets out to carry out a scientific expedition, but unfortunately crash-lands on an alien planet. In order to survive and uncover the mystery of the planet, students solve problems by learning and applying AI knowledge. In the first prototype, the game focuses on learning classical search algorithms: how each algorithm works and what its suitable applications are. To make the learning experience accessible to high school students, the game leverages the math taught in classes before and during high school that is foundational to AI. The game is designed based on the principles of (1) learning through interaction; (2) learning in a familiar context and transfer to a new domain; and (3) focusing on Use (in the Use, Modify and Create phases of CT) (Lytle et al., 2019). We focus on the Use phase as it is the necessary foundation for AI Consumers, Operators, and Developers alike. By focusing on the Use phase, we can tailor our design to provide the understanding of how these algorithms work needed for AI Operators and Developers to later move to the Modify and Create phases, without overwhelming and alienating students less interested in these applications. As a result, the current implementation of the game introduces students to the concepts behind search algorithms, walks them through classical search algorithms, illustrates the strengths and weaknesses of each algorithm, and engages them in a transfer-learning task.

Students are first introduced to the Breadth First Search algorithm through an introductory tutorial task. In this tutorial task, the student must track down their robot companion, named Tachi, who got lost in the crash. The student is provided with a map detailing where Tachi is and must plot a path to it from their current location. The map consists of a series of labeled areas connected by walking paths. This initial problem enables us to introduce the concept of search trees quite easily, as the labeled areas and walking paths on the map readily translate to nodes and edges in a search tree. The spatial nature of this search problem provides a more accessible entry point for students to begin thinking about search and serves as a familiar reference while trying to interpret and understand the novel and abstract tree representation As the student clicks on the areas on the map, the corresponding node is placed in the search tree. This results in the student engaging in Breadth First Search through their own exploration of the map. After several steps of this traversal, the student will be prompted to identify which nodes have already been explored, and which nodes are currently being held in memory. After a path is found, the student is prompted to proceed to the robot's location. Upon finding the robot, the student proceeds with a more difficult puzzle task, in which they must fix the robot's broken circuit board. This is presented in the form of a sliding 8 puzzle (as seen in Fig. 1), with parts of the circuit board being moveable pieces. In order to draw a connection between this task and the previous tutorial task, the student is prompted to try sliding pieces around, and is shown how each new configuration resulting from a move can be represented as a node in a search tree, with the initial configuration being the root. After discovering that Breadth First Search can be applied to this problem, the

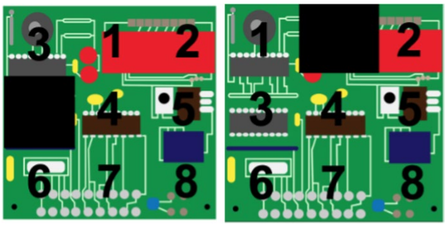

Fig. 1 Representation of a circuit as an 8-Puzzle Game, with initial state on the left and goal state on the right

student then proceeds to walk through the first few steps of the algorithm before the search autocompletes and the solution is found (Fig. 2).

The subsequent module introduces students to the Depth First Search algorithm. In the tutorial task for this algorithm, the student must find a path to a source of materials that they can use to repair their ship. Similar to the Breadth First Search tutorial task, the student is provided with a map detailing both their current location and the location of the materials. Again, this initial engagement with a spatial search space provides an intuitive reference point for further learning. We design this search problem to expose the flaws of Breadth First Search, namely by creating a map that results in a tree that is extremely wide. This exposes that Breadth First Search can result in running out of memory when used on a tree like this. We use this as an opportunity to link AI learning with previous math knowledge. We discuss the time and space complexity, where the student comes to understand them in terms of linear and exponential functions from their math classes. This will help the student determine which algorithm to use in any particular task. The student calculates the number of nodes in memory to be exponential as a function of depth level. The student is prompted to try a different approach to exploring the tree in response to this roadblock. The suggestion to try "picking a direction and continuing until either the goal or a dead end is reached" is given, and by performing this the student performs Depth First Search. When the student is prompted to analyze the number of nodes held in memory when using Depth First Search on this example, it is pointed out that the number of nodes in memory grows roughly linearly with respect to depth level, as opposed to the exponential growth rate of Breadth First Search's memory usage. This analysis grounded in the student's math knowledge forms a basis for future decisions between search algorithms. Upon completing this tutorial task and proceeding to the mineral source on the map, the student finds another terminal to unlock. In

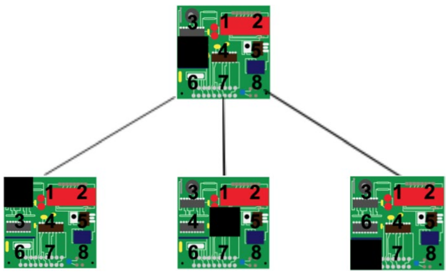

Fig. 2 Search tree representation of the circuit problem up to the first level

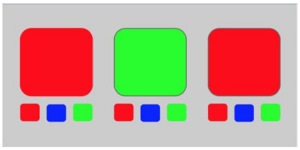

Fig. 3 Representation of a color based code cracking puzzle

order to unlock this terminal, the student must solve a color code puzzle consisting of slots that can each be set to either red, blue, or green (shown in Fig. 3). After Depth First Search is used to solve the puzzle, the student is prompted to calculate what the number of nodes in memory would be if Breadth First Search was used, to determine whether or not it could also be used to solve this puzzle. This problem is designed to show the student that there exist problems that can be feasibly solved by either algorithm.

After solving this puzzle, the student must make their way back to the ship with the materials they have gathered. The student is once again given a map and is tasked with plotting a path to the ship's location. This problem was designed to highlight the flaws of Depth First Search, and to demonstrate the strengths of Greedy Search. Once again, the map results in a tree too wide for Breadth First Search, but the connectivity of areas on the map demonstrates how Depth First Search can get caught in infinite loops in such scenarios. At this point, the student is told that the ship is located near a waterfall, and it is explained how the loudness of the noise from this waterfall can be used as a heuristic to estimate proximity to the ship, demonstrating how heuristics can be used in search. Using Greedy Search to find the shortest path, the student learns that Greedy Search uses minimal memory, avoids getting caught in loops, and can find the shortest path. For the final task, Greedy Search is used to solve a sliding circuit board problem similar to the one in the Breadth First Search puzzle. This problem aims to explore how Greedy Search can be applied to an abstract problem by using a familiar example. The number of out of place tiles is used as a heuristic, one that the student can directly calculate for each node. By doing so, the student learns how to test the effectiveness of a heuristic.

Design Philosophy

Addressing Common Needs

The design of our game-based learning environment outlined in the previous section addresses the common needs of the high school STEM students identified in the cognitive interview study. First, the tutorial task for each algorithm involves using search to find a path to a predetermined location on a map with clearly labeled open areas and walking paths connecting these areas. This arrangement allows for very explicit scaffolding with respect to learning the concept of search trees, an abstract representation inherent to the AI problem of Classical Search. By introducing the concept of search along with this task, it established a solid foundation for understanding the search tree representation before students were prompted to understand how it applied to tasks like sliding puzzles or code cracking. This structure also was designed to support explicit scaffolding with respect to understanding how to parse problems in terms an AI system can operate on. By beginning each algorithm tutorial with a spatial search task using a map that easily lends itself to search tree representations, the student is easily able to conceptualize how a search algorithm solves a problem by operating on an abstract representation that still closely resembles the original framing of the problem. This makes it possible to build upon that understanding for the puzzle tasks that less closely resemble searching through a search tree at first glance. By understanding moving between designated areas via a walking path as traversing between nodes via edges, it becomes possible to transfer that understanding to moving between puzzle configurations by moving pieces. Finally, our design involved explicitly highlighting aspects of mathematical analysis relevant to the algorithms at hand to help students apply their mathematical knowledge to what they were learning. The analyses of time complexity and space complexity were supported by in-the-moment scaffolding, starting with basic counting and comparisons and then building up to analyzing how these quantities grow for specific algorithms as functions of branching factor and tree depth and ultimately comparing these functions to assess the relative strengths and weaknesses of each algorithm.

Ensuring Appropriate Takeaways for All AI Students

Our game-based learning environment was also designed to ensure that each of the three kinds of AI Students came away from the experience with the desired takeaways. For the AI Consumer, they would begin to understand some of the tasks that AI could accomplish and what kinds of information it would need to effectively accomplish this task. In particular, experiencing the difference between brute force search algorithms like Breadth and Depth First Search, and informed search algorithms like Greedy Search gives AI Consumers a glimpse into how more specific information about a task enables more efficient and tailored approaches. This understanding can thus begin to be extended to how high impact applications such as health diagnosis, GPS navigation, and targeted advertising draw their effectiveness from the personal information they have access to. Furthermore, being exposed to simplified versions of real-world applications of AI such as code cracking can begin to give AI Consumers a more concrete understanding of how to make personal decisions to protect themselves from negative applications of these algorithms.

We anticipate that AI Operators would be able to come away with this same level of knowledge, in addition to a deeper understanding of how to differentiate between AI technologies and choose the right one for a particular task. In ARIN-561, students are exposed to the different search algorithms in the context of their advantages and disadvantages. In order to introduce Depth First Search, the student is presented with a problem that Breadth First Search cannot solve due to running out of memory. This prompts a comparison of Depth First Search and Breadth First Search on the basis of memory usage and runtime. Similarly, the introduction of Greedy Search comes in the context of a problem that neither Breadth First Search or Depth First Search can solve. After having walked through the three search algorithms presented and their tradeoffs in terms of space and time complexity, AI Operators will be able to utilize these concepts when framing the tasks towards which they are expected to apply AI technologies. The "puzzle" tasks further build upon students' ability to perform this analysis, by presenting abstract search problems that the student must apply these algorithms to, at times analyzing how different algorithms fare. This directly mirrors the kinds of thinking AI Operators will have to use in making real world decisions about what AI technologies are best suited to a particular problem. This will enable AI Operators to make more informed decisions about the tools they use to accomplish tasks vital to their careers.

AI Developers, in addition to all of the previously mentioned takeaways, should also come away with a deeper understanding of how each of these algorithms works. Having been taken through step by step how each algorithm iterates until its completion, AI Developers get knowledge as to how these search algorithms function in detail. In addition, AI Developers learn not only what time and space complexity are, but introductory ways to calculate and understand it. By introducing the concepts first through counting nodes that were explored or in memory, then tying those numbers to the framework of growth as the depth of the search tree increases, then comparing those growth rates, the design of our game-based learning environment provides a sound basis for understanding asymptotic notation and more rigorous comparisons of space and time complexity. Finally, AI Developers also experience how algorithms like Greedy Search can be tailored to the problem at hand, and how doing so can potentially lead to vastly more efficient task completion.

Narrative Design

As aforementioned, effective narrative design has been suggested to play a significant role in aligning student learning with their engagement in the game-based learning environment. Namely that tight integration between educational content, narrative, and gameplay, along with properly calibrated narrative complexity are two factors that can lead to a positive relationship between learning and engagement in game-based learning environments (Rowe et al., 2011). We incorporated the first factor into our design by building the game's narrative around tasks commonly solved by search algorithms (path finding, solving puzzles, password cracking). After having crash landed on the planet, the student needs to find their robot companion, Tachi, in order to begin fixing their ship, necessitating pathfinding and the introduction of BFS. From that point on, the student only advances the narrative through gameplay that teaches them more about various search algorithms. The second factor calls for a narrative that is detailed enough to motivate the student to solve the problems they are presented with, but not so complex that it distracts the student from learning. Our narrative utilizes a straightforward premise, escape the planet you crash landed on with the help of Tachi, but uses the mysteries surrounding both the planet and the circumstances of your crash to motivate the student to proceed through the gameplay. The inherent mystery of exploring another planet allows us to build the narrative around exploration in a way that piques student curiosity, an important element of promoting engagement in game-based learning environments (Malone, 1981; Provenzo, 1991). The explicit goal of the game, repairing the ship and leaving the planet, as well as the implicit goals of exploring the mysteries of the planet, are all deeply tied to the tasks presented to the student. In this way, we utilize the student's engagement to motivate them in problem solving. The only way to find materials your ship needs, or navigate strange phenomena on the planet, is to utilize an understanding of the Search algorithms presented and use them to approach these situations systemically. Our design drives this engagement even further by developing a relationship between the student and the NPC Tachi. The first AI problem the student must solve is finding Tachi, their robot companion. Then, they must use what they've learned to fix Tachi's broken circuitry. This series of events serves to build companionship between Tachi and the student from the beginning of the game. Throughout the rest of the game, we emphasize the relationship between the student and Tachi through dialogue, engendering the idea that you are two friends, in this situation together and will get through it by working together. Thus, when Tachi helps the student with a particular problem, the student will be driven by their engagement with Tachi as a character to solve it. These factors build a narrative flow that keeps students engaged without distracting from the AI tasks at hand.

Discussion

In this paper, we discussed the design of a game-based learning environment to meet the challenges of teaching AI to high school students. This work was largely based upon the initial findings of a cognitive interview study to uncover how high school STEM students approach AI concepts. An overarching theme emerging from the interviews is that artificial intelligence represents a novel and mysterious problem space for high school aged students. Therefore, one cannot assume facile transfer from grade-level mathematics and computer science concepts to AI problems, even among students with mastery of the underlying concepts. Rather, it is likely that students will need explicit support to recognize and flexibly apply the background knowledge they may have in service of AI problems: there is little evidence from this study that AI can be successfully approached as a near-transfer task in which students can be expected to readily apply knowledge from one context to another. At the same time, we do have evidence that when provided explicit support to incorporate prior knowledge and skills into an AI learning experience, students are adept at leveraging this knowledge to solve AI problems. This suggests that AI may provide a powerful vehicle to deepen mathematical and computational thinking as students are compelled to expand beyond a school-bound understanding of mathematics as they apply it to solve compelling AI problems.

Similarly, our findings about student difficulties with common AI representations like search trees are worth considering alongside findings about student difficulty with mathematical representations. This challenge speaks to the role of computational thinking for successful engagement with AI problems, which we observed not only in students' difficulties with the abstractions central to AI problem solving approaches, but also in their somewhat tenuous grasp of abstractions inherent in common mathematical representations such as tables and graphs. This finding dovetails with longstanding research about student difficulty understanding the mathematical relationships represented in graphs and tabular data (see, e.g., Curcio, 1987). Whereas it is unsurprising that students 13–17 years old are unfamiliar with search tree representations or data frames, such abstractions are critical in understanding how information may be structured in ways that enable AI systems to solve problems yet may present stumbling blocks without explicit support.

Our game-based learning environment was designed to effectively address these needs for scaffolding with respect to abstract and mathematical representations, as well as understanding how AI systems operate on problems. By building explicit scaffolding to teach how to build search trees and how search algorithms operate on them, as well as supporting mathematical analysis of the search space and time complexities, we have centered addressing these needs in our design. At the same time, our learning environment was crafted to ensure that AI Consumers, Operators, and Developers would all be able to come away with understanding valuable to their relationship with AI. AI Consumers gain exposure to real world applications of Classical Search and the information needed to accomplish them, AI Operators learn how to choose the right AI technology for a particular Classical Search task, and AI Developers come away with a step by step understanding of how these algorithms work. Finally, the narrative of our game-based learning environment was designed to encourage a positive relationship between engagement and learning through a tight integration between educational content, narrative, and gameplay, as well as properly calibrated narrative complexity based on recommendations from the success of educational games like Crystal Island (Rowe et al., 2011). Our narrative, based around escaping an alien planet after a crash landing, is built around a straightforward premise that enables both embedded mystery to spark curiosity and tasks that align game progression with AI knowledge advancement.

Limitations

One limitation of our proposed design is that it is largely informed by an exploratory study with limited sample size. Furthermore, this exploratory study was conducted at a private school with a strong focus on STEM education. We suspect that students in schools with fewer resources and/or less of a focus on STEM/Computational Thinking, may have different or additional needs that our design has not as yet addressed. This also holds true for students from different grade levels within K-12 education, or students from demographics not sufficiently represented in our sample. In addition, this design has not yet been evaluated in any studies with students to fully evaluate whether it does indeed address these needs and lead to the desired takeaways for all categories of AI students.

Future Work

To address the aforementioned limitations, we plan to run studies with K-12 students to evaluate the effectiveness of our design. We will begin with a study at the same school where we performed our initial exploratory study, then aim to branch out to various schools around the country. More specifically, we aim to both revisit our cognitive interview work and evaluate the effectiveness of our design in more varied K-12 environments (public schools, different grade levels), and aim to sample students from demographics not sufficiently represented by our exploratory study.

Acknowledgements This research was supported by the National Science Foundation under Grant 1842385. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Funding Open access funding provided by SCELC, Statewide California Electronic Library Consortium

Declarations

Competing Interests The authors have no relevant financial or non-financial interests to disclose.

References

Ali, S., Payne, B. H., Williams, R., Park, H. W., & Breazeal, C. (2019). Constructionism, ethics, and creativity: Developing primary and middle school artificial intelligence education. In *International Workshop on Education in Artificial Intelligence K-12 (EDUAI'19).*

Basu, S., Biswas, G., Sengupta, P., Dickes, A., Kinnebrew, J. S., & Clark, D. (2016). Identifying middle school students' challenges in computational thinking-based science learning. *Research and practice in technology enhanced learning, 11*(1), 1–35.

Batanero, C., Godino, J. D., Vallecillos, A., Green, D. E., & Holmes, P. (1994). Errors and difficulties in understanding elementary statistical concepts. *International Journal of Mathematical Education in Science and Technology, 25*(4), 527–547.

Chen, Y., & Tang, X. (2018). *Fundamentals of artificial intelligence for high schools*. East China Normal University Press.

Curcio, F. R. (1987). Comprehension of mathematical relationships expressed in graphs. *Journal for Research in Mathematics Education, 18*(5), 382–393. https://doi.org/10.5951/jresematheduc.18.5.0382

Du Boulay, B. (1986). Some difficulties of learning to program. *Journal of Educational Computing Research, 2*(1), 57–73.

Friel, S. N., Curcio, F. R., & Bright, G. W. (2001). Making sense of graphs: Critical factors influencing comprehension and instructional implications. *Journal for Research in Mathematics Education, 32*(2), 124. https://doi.org/10.2307/749671

Gardner-McCune, C., Touretzky, D., Martin, F., & Seehorn, D. (2019). 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, SIGCSE '19* (pp. 1244–1244). ACM.

Gee, J. P. (2003). *What video games have to teach us about learning and literacy*. Palgrave Macmillan.

Greenwald, E., Leitner, M., & Wang, N. (2021, May). Learning artificial intelligence: insights into how youth encounter and build understanding of AI concepts. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35, No. 17, pp. 15526–15533) with a pre-print retrieved May 4, 2022 at https://www.aaai.org/AAAI21Papers/EAAI-78.GreenwaldE.pdf

Guszcza, J., Lewis, H., & Evans-Greenwood, P. (2017). Cognitive collaboration: why humans and computers think better together. *Deloitte Review, 20*, 8–29.

Heintz, F., Mannila, L., Nygårds, K., Parnes, P., & Regnell, B. (2015, September). Computing at school in Sweden–experiences from introducing computer science within existing subjects. In *International Conference on Informatics in Schools: Situation, Evolution, and Perspectives* (pp. 118–130). Springer.

Kandlhofer, M., Steinbauer, G., Lasnig, J. P., Baumann, W., Plomer, S., Ballagi, A., & Alfoldi, I. (2019). *Enabling the Creation of Intelligent Things: Bringing Artificial Intelligence and Robotics to Schools. 2019 IEEE Frontiers in Education Conference (FIE)*. https://doi.org/10.1109/fie43999.2019.9028537

Lee, S., Mott, B., Ottenbreit-Leftwich, A., Scribner, A., Taylor, S., Park, K., … Lester, J. (2021). AI-Infused Collaborative Inquiry in Upper Elementary School: A Game-Based Learning Approach. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35, No. 17, pp. 15591–15599).

Lytle, N., Cateté, V., Boulden, D., Dong, Y., Houchins, J., Milliken, A., Isvik, A., Bounajim, D., Wiebe, E., & Barnes, T. (2019). Use, Modify, Create: Comparing Computational Thinking Lesson Progressions for STEM Classes. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '19)*. Association for Computing Machinery, 395–401. https://doi.org/10.1145/3304221.3319786

Malone. (1981). Toward a theory of intrinsically motivating instruction. *Cognitive Science, 4*, 333–369.

Makridakis, S. (2017). The forthcoming Artificial Intelligence (AI) revolution: its impact on society and firms. *Futures, 90*, 46–60.

Nadikattu, R. R. (2016). The emerging role of artificial intelligence in modern society. *IJCRT - International Journal of Creative Research Thoughts (IJCRT), 4*(4), 906–911. https://doi.org/10.1729/journal.24110

Pimentel, C. S., Queiroz, R. L., Lima, P. M. V., & Sampaio, F. F. (2018). Projeto Frankie: uma proposta para o ensino de Inteligência Artificial na Educação Básica. *Nuevas Ideas En Informática Educativa, 14,* 493–498. http://www.tise.cl/Volumen14/TISE2018/493.pdf

Provenzo, E. (1991). *Video kids: making sense of Nintendo /*. Harvard University Press,.

Rich, K. M., Yadav, A., & Zhu, M. (2019). Abstraction in students' mathematics strategies: productive starting points for introducing CT concepts. *Journal of Computers in Mathematics and Science Teaching, 38*(3), 267–298.

Ritter, S., Aglio, J., Stetzer, R., Wilson, G., Navta, N., & Lippl, C. (2019). *Teaching AI as computational thinking for middle school students*. AIED Workshop.

Rodríguez-García, J., Moreno-León, J., Román-González, M., & Robles, G. (2021). Evaluation of an online intervention to teach artificial intelligence with learningML to 10-16-year-old students. *Technical Symposium on Computer Science Education.* https://doi.org/10.1145/3408877.3432393

Rowe, J. E., Shores, L. R., Mott, B. W., & Lester, J. C. (2011). Integrating learning, problem solving, and engagement in narrative-centered learning environments. *Artificial Intelligence in Education, 21*(1), 115–133. https://doi.org/10.3233/jai-2011-019

Russell, S. J., & Nerving, P. (2016). *Artificial intelligence: A modern approach,* Malaysia.

Salen, K. (2007). Gaming literacies: a game design study in action. *Journal of Educational Multimedia and Hypermedia, 16*(3), 301–322.

Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019). Blue Sky: Envisioning AI for K-12: What should every child know about AI? In *Proceedings of AAAI '19*, Jan 27 – February 1, Honolulu, HI. AAAI Press.

Wang, N., & Johnson, M. (2019). *AI education for K-12: connecting AI concepts to the High School Math Curriculum*. IJCAI EduAI workshop.

Wouters, P., van Nimwegen, C., van Oostendorp, H., & van der Spek, E. D. (2013). A meta-analysis of the cognitive and motivational effects of serious games. *Journal of Educational Psychology, 105*(2), 249–265. https://doi.org/10.1037/a0031311

Xiong, Y., Wang, J., & Huang, J. (2018). *Textbook series on artificial intelligence for elementary and middle schools*. East China Normal University Press.

Yeh, T., Dalton, B., & Haberl, E. (2019). Integrating AI ethics into robotics learning experiences. In *Proceedings of the 20th International Conference on Artificial Intelligence in Education - Workshop on K12 AI education*.