

Learning Artificial Intelligence: Insights into How Youth Encounter and Build Understanding of AI Concepts

Eric Greenwald,¹ Maxyn Leitner,² Ning Wang³

University of California, Berkeley, Lawrence Hall of Science¹
University of Southern California, Institute for Creative Technologies^{2,3}
eric.greenwald@berkeley.edu,¹ mjohnson@ict.usc.edu,² nwang@ict.usc.edu³

Abstract

Artificial Intelligence's impact on society is increasingly pervasive. While innovative educational programs are being developed, there has been little understanding of how students, especially pre-college aged students, construct an understanding of and gain practice with core ideas about AI or what concepts are most appropriate for what age-levels. In this paper, we discuss a cognitive interview study with middle school and high school students to better understand how students learn AI concepts. We aim to shed light on questions including: what is the range of background knowledge and experiences students are able to apply in encountering AI concepts; what concepts are most readily accessible and which are more challenging; what misconceptions do students bring to bear on AI problems; and how to help students approach AI concepts by leveraging related concepts (such as mathematical and computational thinking). Results from the exploratory study have the potential to provide important insights into AI learning for pre-college youth. These initial findings can inform further investigations to ground the design of learning and assessment in evidence-based learning progressions and grade-level performance expectations.

Introduction

Math and computational thinking are two important skills for the nation's future workforce and are foundational to the field of artificial intelligence (AI). Mathematics is foundational to computer science, a field that combines mathematics, engineering, and science (Denning, 2009). Computational thinking (CT), encompassing a broad range of mental tools and concepts from computer science, helps people solve problems (e.g., diagnosing disease), design systems (e.g., self-driving cars), understand human behavior (e.g., speech recognition), and engage computers in

automating a wide range of intellectual processes (NRC, 2010). 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, Lewis, & Evans-Greenwood, 2017). Thus, proficiency in the language of AI is key to a data-capable 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, while AI's impact on society is increasingly pervasive, and innovative educational opportunities are being rapidly developed, there has been precious 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.

In this paper, we discuss a cognitive interview study to better understand how students learn AI concepts. The study is part of a project that aims to develop game-based learning environments to help students build AI competency, apply

math knowledge and develop CT skills. We conducted surveys and cognitive interviews to answer questions including: what is the range of background knowledge and experiences students are able to apply in encountering AI concepts; what concepts are most readily accessible and which are more challenging; what related concepts (e.g., mathematical and computational thinking) do students engage with and or need to leverage in order to access AI concepts? While the initial work presented here is exploratory, the survey data and in-depth interviews with students have the potential to provide important insights into AI learning for pre-college youth. The preliminary results can inform future studies that systematically examine these initial findings and support the advancement of the field's understanding of student knowledge construction of AI concepts and thereby better support the design of instruction and assessment aimed at pre-college youth.

Related Works

While 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, with the exception of China (Xiong, Wang, & Huang, 2018; Chen & Tang, 2018), which has already developed a series of seven AI textbooks for elementary, middle, and high schools. Sweden has also developed AI courses to educate its citizens, including school-age youth, about AI (Heintz et al., 2015).

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.² This includes a curriculum for teaching ethics of AI to middle school students.³

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).

The work presented here aims to uncover how K-12, particularly high school students, approach AI concepts, what obstacles they face, and how to guide them through the obstacles. The work builds upon previous investigations into linking AI to K-12 math curriculum to identify AI concepts suitable for high school students. (Wang & Johnson, 2019), as well as work investigating the learning of computational thinking (Lee, et al., 2011; Rich, et al., 2019) and seminal research into comprehension of mathematical representations (e.g., Curcio, 1987; Friel, Curcio & Bright, 2001) and statistics (e.g., Batanero, Godino, Vallecillos, Green, & Holmes, 1994).

Methodology

We conducted a study using cognitive interviews with middle school and high school students using five AI problems to gain an understanding of what type of AI concepts and difficulty is suitable for the high-school population, what are the challenges they face, and what pedagogical approach can be applied to guide students in AI problem-solving.

Sample

We recruited 8 participants from a private school located in the United States. The participants consisted of 4 high school students and 4 middle school students, ranging in age from 12 to 17 years old. The inclusion of the middle school students is based on suggestions from the teachers at the school, who believed the AI problems might be accessible to advanced middle school students after reviewing the AI problems used in the study and the math knowledge required. The sample was drawn from and reflected a student population that is racially and ethnically diverse and generally higher-resourced than that of the average public school. Especially given persistent inequities in computer science academic and career pathways, as well as broader questions of access and privilege across society, the sample presents limitations for generalizability that will be important to address in future studies.

Measures

As students began the cognitive interview, they were asked background information such as grade and age. Then, students were asked to rate their interest in AI, their understanding of AI, and their confidence in math class, each on a Likert scale of 1 to 5.

The cognitive interview employed a semi-structured protocol, with a series of prompts to elicit student thinking as they a) initially encountered each problem (e.g., “Does anything about the problem seem familiar to you?”, “What do you think a ‘good answer’ might look like?”); b)

² aieducation.mit.edu/

³ www.media.mit.edu/projects/ai-ethics-for-middle-school/overview/

attempted to solve the problem (e.g., “I see that you [did X], tell me about your thinking; and, how did you decide to do X?”); and c) after they had settled on a solution (e.g., “Were there any moments when the problem became more clear for you or where you noticed that you had a better understanding of how to solve it?”). In-the-moment scaffolding was provided throughout each interview to: a) enable students to reveal thinking across each step of the solution; b) surface and test emerging ideas about why a student might be stuck; and c) disambiguate between superficial challenges, such as unfamiliar vocabulary and calculation errors, and more conceptual difficulties.

AI Material

To determine the AI concepts to teach, we reviewed the AI curriculum from the most popular AI textbook for higher education (Russell & Novig, 2016) and categorized the AI topics into four main fields (Search, Knowledge Representation and Planning, Probabilistic Reasoning, and Machine Learning). Based on the underlying math knowledge required, we then narrowed down the sets of AI concepts based on whether the relevant math knowledge is covered by the high school curriculum. 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, we selected five AI concepts – Search, Bayesian Networks, Decision Trees, Clustering, and Linear Regression. Using the classic problems for the selected AI concepts as examples, we designed an AI problem for each of the AI concepts. Each problem consisted of a series of questions that help students develop the solution step by step. The problems were designed to be difficult to complete, with the difficulty increasing at each step toward the solution.

For Classical Search, students were presented with an 8-Puzzle Game (shown in Figure 1). In the game, the students were given a 3x3 board with initial configuration of 8 sliding pieces, numbered 1-8, and were asked to slide the pieces around to reach a target configuration of the board. Students were asked to complete the puzzle and estimate the size of the search tree.




Fig. 1

Representations of the 8-Puzzle Game in its initial and completed states.

For Bayesian Networks, students were presented with a problem to estimate the likelihood that a patient has a cavity. Students were given a problem description with joint probability of events (e.g., “A dentist often says that 10.8% of patients have a cavity and a toothache, and the dental hook catches the tooth.”). They were asked to organize the probability in a joint-probability table (Figure 2), and answer questions about the likelihood of various events by reading the table. The students were then guided to organize the variables in a Bayesian Network.

	Toothache		No Toothache	
	Catch	Not Catch	Catch	Not Catch
Cavity	0.108	0.012	0.072	0.008
No Cavity	0.016	0.064	0.144	0.576

Fig. 2

Tabular data as presented in the Cavity problem.

For Decision Trees, students were presented with a problem where they had to help a friend decide whether or not they would wait to get a seat at a restaurant or not. A table containing data about past experiences detailing whether or not they would wait for a table at restaurants was included. The data table contained features such as the wait time, the type of restaurant, whether it’s raining etc. The students were asked to think of what questions they could ask their friend to help them come to a decision.

For Clustering, students were given a common clustering problem of identifying flower species. Students were given a table of sample data on Iris flowers with features such as petal width, petal length, sepal width, sepal length, and flower species. Students were presented measurements of a new flower and asked to identify the species of the new flower.

For Linear Regression, we designed a problem based on a tutorial on “Exploring bivariate numerical data” from Khan Academy.⁴ Students were presented with a table containing past data on time spent on the phone and the battery life remaining (Figure 3). Students were asked to use the data to predict the percentage of battery life left in a mobile phone after certain hours of use.

Time spent on phone (hours)	1	2	3.5	4	6	7	8	9
Battery life remaining (hours)	8	7	7	5.5	5	3.5	2.5	2.5

Fig. 3

Table of time spent on a phone and corresponding battery life remaining.

Procedure

The cognitive interviews were conducted by a member of the research team with expertise in the learning sciences and mathematics instruction. During the study, the interviewer first asked students about their background information using questions from the Measures section. Then the interviewer presented the AI problems printed on paper to the students. Students were given blank paper and pencils to work on the problems. Students were given blank paper and pencils to work on the problems. Students were encouraged to think aloud as they solved the problems. The interviewer asked questions and guided students through each problem. Every student was presented with all five problems and worked toward a solution (not necessarily completed) for 2-3 problems during the interview. Problem selection included input from the students interviewed and was ultimately determined by the researcher conducting the interview to ensure a balance of data across each problem type. How much of each problem students completed depended on the time the students spent on the problem and whether the interviewer deemed suitable to continue, based primarily on student progress and indications of student frustration levels. Table 1 describes the number of students who engaged with each problem type. Each study session lasted about 1 hour. Top-down videos of the paper the student used to work on the AI problems were recorded.

Problem Type	Responses
Search	5
Bayesian Networks	4
Decision Trees	5
Clustering	4
Linear Regression	4

Table 1

Distribution of problem responses across participants.

Analysis

Video files for each interview were uploaded to Dedoose (2018), a mixed methods data analysis tool. To analyze the interview data, we first created excerpts of each student's work on each problem, and applied codes directly to the video excerpts according to the problem type. We then completed two passes through the data. With the first pass we viewed each interview in sequence, generating and applying a set of broadly applicable codes (e.g., “challenges” to denote moments of student difficulty with a problem, and “prior knowledge” to denote evidence of students applying prior mathematical knowledge to a problem). Then, we compiled the excerpts by problem type and viewed the variety of student responses on each problem together. In this second pass through the data, we drew from the principles of grounded theory (Glaser & Strauss, 1967; Glaser, 1992), to iteratively introduce new codes as themes emerged (e.g., “problem identification,” and “central tendency”), which we then added to and revised through successive passes through the data, continually comparing the emergent codes against the data. These initial codes were revised for consistency and tractability, then applied systematically across the data. While we report on findings that have emerged from these initial passes through the data, additional analyses are ongoing.

Initial Findings

Analysis of the three Likert scale questions for the students interviewed (asked at the outset of the interview) provide useful context for interpreting findings. Likert items can be difficult to compare across respondents, owing to variation in student propensity to anchor responses at the high end or low end of the scale. However, comparisons within each student's responses to the three questions offer a window into a student's relative confidence and interest in AI and math. As revealed in the table below (Table 2), nearly all

⁴ khanacademy.org/math/statistics-probability/describing-relationships-quantitative-data/introduction-to-trend-lines/a/equations-of-trend-lines-phone-data

students expressed high levels of interest in AI, and rated their interest in AI higher than their confidence in understanding AI. At the same time students generally felt more confident with math than AI, although only slightly.

Student Age	Math Confidence	AI Confidence	AI Interest
12	2	3	4
13	3	3	5
13	3	2	3
13	4	2	4
15	4	3	4
17	4	1	2
17	3	2	5
17	2	2	3

Table 2

Student age, confidence in math and AI, and interest in AI.

With this context about the student sample in mind, analysis of the student cognitive interview data has shed new light on how 12-17 year old students encounter and construct knowledge related to artificial intelligence concepts. We organize these early findings into five themes.

Students needed support to leverage and apply mathematical concepts that underlie AI problems

Even in cases where students demonstrated competency with the necessary mathematical skills, they often struggled to identify connections and/or make use of those skills until explicitly prompted. This was particularly evident for the application of statistics and probability (e.g., drawing on probability concepts in approaching Bayesian Networks and regression modeling for the linear regression algorithm). For example, when students were provided data about cell phone battery usage and asked to use the data to predict the percentage of battery life left in a mobile phone, each student initially scanned the values in the data and offered an estimate based on the final 1 or 2 data points provided (see Figure 3). Even after prompting, (such as, “can you think of a way to graph these data to get a more precise answer?”) students struggled to do so until the idea of a line of best fit was explicitly introduced by the researcher conducting the interview. Once introduced, however, students generally recognized the method and were able to apply it to the problem. This theme was also evident in the use of graphical or tabular data representations. For example, the data used for the cavity problem and for the restaurant decision tree problem were both presented in tabular displays. The nested nature of the data table for the cavity problem (see Figure 2), in which various

combinations of conditions were represented together, was particularly challenging for students.

For nearly all students, these data tables were initially difficult to interpret as computational artifacts, with students struggling to make meaning of the relationship between values within a row or to connect column headers to variable descriptions. This led to difficulty in making independent progress on the problems until students were supported to attend to the structural features of the data as represented in the tables.

However, for most students, in-the-moment scaffolding was effective in enabling them to recognize and apply their mathematical knowledge for AI problems. Once the connection between the AI problem and the underlying math was made explicit, students were able to engage productively. For example, in the linear regression problem, all but one student was readily able to construct a graph and produce a reasonable line of best fit to the data after this connection to the math was explicitly introduced. Similarly, explicit scaffolding about the relationships among rows, columns, and cells within the data tables was largely successful in activating students’ mathematical and computational thinking to use the tables to solve the AI problem. This suggests that students will likely need support in identifying when and how the math concepts they may be familiar with in their mathematics classroom can be applied to AI systems. More promisingly, it also suggests that once this background knowledge is activated, students can leverage it to productively engage with AI problems.

Students found difficulty with the abstract representations characteristic in AI problems

Across all interviews, students needed explicit scaffolding in understanding how to interpret and construct a search tree (e.g., in order to interrogate the fitness of different search algorithms for different search problems). Thus, even after the interviewer scaffolding enabled students to construct a search tree from the slider puzzle, students struggled to make use of the search tree representation to consider the depth and breadth of a problem space (both for the search trees they constructed and for pre-constructed exemplars). As an illustration of this, with scaffolding, all students were able to build a tree from a node to branches to new nodes, yet only one student was able to recognize the salience of tree abstractions such as branching factor and tree depth (albeit using colloquial language) to estimate the relative complexity of a problem. While by no means definitive, there were some indications that student difficulty to use, modify, and create search tree or Bayesian Network representations interacted with students’ challenges applying concepts of probability to the AI problem space (i.e., in recognizing the universe of possibilities that systematically determine the construction of tree nodes). With ongoing analyses to better understand this relationship between difficulty with representations and difficulty with the concepts being represented in the case of search trees,

initial findings suggest the importance of supporting students to understand the computational features of tree representations--in essence to learn how to “read” a tree.

Students first draw on their own experiences with a problem’s context when approaching AI problems

In a variation of the human-interpreter problem (Spohrer & Soloway, 1986) common to computer science, several students began their attempt at solving a problem by referring to prior experience with the problem space (e.g., cell phone batteries and restaurant dining) in an attempt to reason through to an answer. For example, in the decision-tree problem, students were reluctant to systematically examine the data provided, and instead began by considering what they would do or what they did in the past (e.g., “if it’s an expensive dinner, i think they’d be more willing to wait.”). That is to say, rather than attending to how features of the data could be operated on to produce a decision, students first tried to figure out the motivations for the agents in the problem (i.e., the diners). Thus, students’ initial answers were less related to the available data, and more related to how the students imagined the diners would behave. We saw similar approaches in the linear regression problem, where students’ initial predictions were often anchored in their own experience with cell phone battery life, rather than the presented data.

Students are unfamiliar with parsing the world in terms an AI system can operate on

A challenge for all students interviewed, even those with advanced mathematical skills, was recognizing how a problem in the world could be made amenable to the computational power of AI. That is to say, students needed support in conceiving a problem space in a way that would enable an AI system to solve it. Thus while some students in the study volunteered ways a computer program might be able to implement an AI solution once identified, the initial step of reconceiving a problem as an AI problem was elusive: the broad strategies AI systems leverage to make predictions or to find a solution from an array of possibilities were unknown to students and thus unavailable resources in their mental models of the problem space. For example, for most students the idea that a problem like an 8-puzzle could be reconceived as a search for a solution was difficult to make use of as the mechanics of AI “search” was largely a black box for which they had little in the way of working theories. Accordingly, there was little evidence that students conceived of search as a potentially systematic process, or that a search space had dimensions that could shed light on the computational difficulty of the search problem. We saw similar evidence in the way students encountered problems for which AI solutions depended on appreciating the relative value of information and the idea of entropy, particularly as represented in the decision tree problem.

Clustering may serve as a productive on-ramp to learning AI concepts

From the set of AI concepts under investigation, students most consistently demonstrated an intuitive grasp of clustering. Even for the middle-school aged students in the study, the concept of centroids and clustering were readily grasped and applied to AI classification problems. In contrast with other problems (as discussed above) in which students struggled to apply mathematical concepts in AI, students needed little scaffolding to apply notions of mean and central tendency to these problems. For example, after students were supported to understand the tabular data needed for classification, each was asked, “how might you go about identifying the unknown species?” In response, students began by looking across the values for different features (e.g., “petal width” and “petal length”) and indicated that the unknown species was closest to the measurements of one of the candidate flowers. When pressed about how they came to this conclusion, students typically reported that they estimated the average of the values for each feature, either using mathematical language (average or mean) or more colloquial expressions of central tendency (e.g. describing the range of values: “I saw that [the petal lengths for a known species] were between 1.9 and 2.5”).

Also interesting was that these notions of central tendency were fluently leveraged by students to identify which features were the most useful in the classification effort. For example, when asked why a student ignored a feature for their flower identification, the student suggested that it was because of “how spread out it is” compared to the values for other features. This attention to the relative usefulness of a measurement meant that students typically transitioned easily to aspects of the problem involving an understanding of information value. This trend may represent possible counter-evidence to our finding (reported above) about students’ difficulties in parsing a problem: for this problem set, students were readily able to assess the relative value of different information about an unknown flower species (e.g., through the prompt, “if you could only get information about one feature of the flower, which would you want?”). While additional research is needed, we speculate that students’ greater familiarity with identifying unknown species (e.g., from biology class) may have positioned them to more easily recognize that some information is better than others for classification. There was also some indication within the interviews that this conceptual grasp of central tendency was able to bootstrap students’ developing understanding of the distance formula. While all but one student expressed unfamiliarity with the distance formula, few students showed difficulty applying it to clustering problems once introduced.

Discussion

In this paper, we discussed the initial findings of a cognitive interview study to uncover how high school 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, such as tabular data. 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). While 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.

A related theme is that the students we interviewed, many of whom had extensive programming experience and high math competency, were unfamiliar with the strategies designers of AI systems use to represent the world and solve problems. While we were not expecting students to have a technical understanding of AI approaches, we were somewhat surprised at how much of a black box AI systems were to students, even as they recognized the myriad places such systems were employed. Interviews suggest a related need to support students toward a more generalized understanding of how AI systems can be applied to

problems, and how problems can be reimaged to be solvable by AI systems. This finding adds weight to efforts aimed at promoting "explainable AI" (Gunning & Aha, 2019) that makes the decision-making of AI algorithms transparent to users. Through its transparency, explainable AI can create opportunities to make AI concepts accessible, in part by supporting youth in developing working theories about how such systems function.

Acknowledgments

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.

References

- Batanero, C., Godino, J. D., Vallecillos, A., Green, D. E., and 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. and Tang, X. (2018). *Fundamentals of Artificial Intelligence for High Schools*. East China Normal University Press, 2018.
- Curcio, F. R. (1987). Comprehension of mathematical relationships expressed in graphs. *Journal for Research in Mathematics Education*, 382-393.
- Dedoose (2018). Version 8.0.35, web application for managing, analyzing, and presenting qualitative and mixed method research data. Los Angeles, CA: SocioCultural Research Consultants, LLC www.dedoose.com.
- Denning, P. J. (2009). The Profession of IT Beyond Computational Thinking. *Communications of the ACM*, 52(6), 28-30.
- Friel, S. N., Curcio, F. R., and Bright, G. W. (2001). Making Sense of Graphs: Critical Factors Influencing Comprehension and Instructional Implications. *Journal for Research in Mathematics Education*, 124-158.
- Gardner-McCune, C., Touretzky, D., Martin, F., and 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, pages 1244-1244, New York, NY, USA. ACM.
- Glaser, B. (1992). *Basics of Grounded Theory Analysis*. Mill Valley, CA: Sociology Press.
- Glaser, B. G., & Strauss, A. L. (1967). *Discovery of Grounded Theory: Strategies for Qualitative Research*, Routledge.
- Gunning, D., & Aha, D. W. (2019). DARPA's Explainable Artificial Intelligence Program. *AI Magazine*, 40(2), 44-58.
- 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., Nygård, 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, Cham.

Holzinger, A., Plass, M., Holzinger, K., Crişan, G. C., Pintea, C. M., & Palade, V. (2016, August). Towards interactive Machine Learning (iML): applying ant colony algorithms to solve the traveling salesman problem with the human-in-the-loop approach. In International Conference on Availability, Reliability, and Security (pp. 81-95). Springer, Cham.

Lee, I., Martin, F., Denner, J., Coulter, B., Allan, W., Erickson, J., ... and Werner, L. (2011). Computational Thinking for Youth in Practice. *ACM Inroads*, 2(1), 32-37.

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. *International Journal of Creative Research Thoughts*.

National Research Council. (2010). *Report of a Workshop on the Scope and Nature of Computational Thinking*. National Academies 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.

Russell, S. J., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*. Malaysia.

Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K. R. (Eds.). (2019). Explainable AI: Interpreting, Explaining and Visualizing Deep Learning (Vol. 11700). Springer Nature.

Shannon, C. E. (1956). A Universal Turing Machine with Two Internal States. *Automata Studies*, 34, 157-165.

Spohrer, J. C., & Soloway, E. (1986). Novice mistakes: Are the folk wisdoms correct?. *Communications of the ACM*, 29(7), 624-632.

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, Palo Alto, California USA

Wang, D., Yang, Q., Abdul, A., & Lim, B. Y. (2019, May). Designing Theory-Driven User-Centric Explainable AI. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1-15).

Wang, N., Johnson, M., (2019) AI Education for K-12: Connecting AI Concepts to the High School Math Curriculum. IJCAI EduAI workshop.

Weintrop, D., Beheshti, E., Horn, M., Orton, K., Jona, K., Trouille, L., & Wilensky, U. (2016). Defining Computational Thinking for Mathematics and Science Classrooms. *Journal of Science Education and Technology*, 25(1), 127-147.

Wing, J. M. (2008). Computational Thinking and Thinking about Computing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1881), 3717-3725.

Xiong, Y., Wang, J., and Huang, J. (2018). *Textbook Series on Artificial Intelligence for Elementary and Middle Schools*. East China Normal University Press.