

AI in Computer Science Education: Tool, Subdomain, and Wildcard

Julie M. Smith*,

* Institute for Advancing Computing Education, Peoria, IL, USA
julie@csedresearch.org

Abstract—The future of AI will be determined in part by how its developers are educated. Thus, how computer science (CS) education incorporates instruction in various aspects of AI will have a substantial impact on AI’s evolution. Understanding how and what CS educators think about AI education is, therefore, an important piece of the landscape in anticipating – and shaping – the future of AI.

However, little is known about how educators perceive the role of AI education in CS education, and there is no consensus yet regarding what AI topics should be taught to all students. This paper helps to fill that gap by presenting a qualitative analysis of data collected from high school CS instructors, higher education CS faculty, and those working in the tech industry as they reflected on their priorities for high school CS instruction and on anticipated changes in high school, college, and workplace CS. We conclude with recommendations for the CS education research community around AI in K-12, particularly at the high school level.

Keywords—computer science education, artificial intelligence

I. INTRODUCTION AND BACKGROUND

Understanding how and what CS educators think about AI education is a significant component of anticipating – and shaping – the future of AI, but there is no widespread agreement as to what AI content should be taught to students (1). Thus, we seek to answer the research question: How do CS educators conceptualize AI education?

The recent expansion of both AI technology as well as public awareness of and interest in AI has led to a variety of opportunities and challenges in the educational domain. In a recent report, UNESCO outlined six challenges related to AI in education: (1) developing comprehensive policies on AI, (2) ensuring equity and mitigating bias, (3) providing teacher professional development, (4) collecting high-quality data for AI systems, (5) implementing a robust research agenda on AI in education, and (6) promoting ethical data use (2). Several groups – including Code.org, the International Society for Technology in Education, and the World Economic Forum – recently released a toolkit designed to assist school systems in developing policies related to the use of AI in schools (3). It catalogues both potential benefits (e.g., the ability to differentiate instruction) and challenges (e.g., perpetuating biases) of AI use in educational contexts. A prominent concern is the

ability of students to use new AI tools, such as ChatGPT, in ways that their instructors consider plagiarism. The advent of plagiarism detectors followed quickly on the heels of public access to ChatGPT, but concerns around their accuracy and fairness exist, such as the fact that current detectors are more likely to incorrectly identify as AI-generated the work of writers whose first language is not English (4).

In addition to the opportunities and challenges that AI presents to all educational settings, there are issues distinct to CS education. Some school systems – such as those in California (5) and Gwinnett County, Georgia, (6) – have already articulated positions on AI instruction. Efforts to teach students about AI are relatively new, especially at the introductory level, although these efforts are rapidly multiplying. Some research has shown that non-CS majors in their first year of college can have positive experiences learning about AI, including some of its technical aspects (7). However, the researchers note that their initial expectations for how much content could be covered were unrealistic, requiring them to scale back some components of the course. Other research has shown that AI topics, including machine learning, can be successfully taught to students in high school and even in middle school (8; 9). Another initiative used AI-focused CS instruction in a summer program for high school girls, resulting in technical understanding as well as in increased confidence and interest in CS (10). As AI education extends into even younger grades, the focus is often on ‘big ideas’ such as the concept that computers can learn from data (11; 12) or the development of ‘AI thinking’ akin to computational thinking (13).

A framework suggested for K-12 AI education proposes three areas of instruction: AI concepts, AI applications, and AI ethics/safety (14). A similar framework advocates that K-12 students learn about AI from a technological perspective, a user-oriented perspective, and a socio-cultural perspective (1). Both of these frameworks emphasize the necessity for all students to gain a basic understanding of AI; they also promote instruction in the ethical and social implications of AI.

II. METHODOLOGY

As part of a project with the goals of (1) articulating what CS content is essential for high school students and (2) defining pathways for high school CS study,

This material is based upon work supported by the U.S. National Science Foundation under Grant No. 2311746. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

feedback was solicited from a variety of relevant parties. Participants were recruited from the CS education practitioner and researcher communities; of the 298 individuals who expressed interest, 21 were included in the focus groups, 38 in the in-person meetings, and 134 provided asynchronous feedback. Participants were selected based on a semi-random process that prioritized diversity across several dimensions, including geographic, expertise, role, and personal demographics.

We hosted three focus groups (one each of high school CS teachers, higher education CS faculty, and those with CS industry experience), each meeting virtually for three one-hour sessions. Topics for these focus groups varied according to the role of the participants and included topics such as priorities for high school CS, anticipated industry trends, possible pathways for high school CS, and anticipated changes in college-level CS instruction. At a two-day in-person meeting held in November 2023, participants had several opportunities to provide input on what content they prioritized in a foundational high school CS course. We also solicited asynchronous feedback, asking participants to vote on their top priorities for a foundational high school CS course.

The data analyzed in this project includes focus group transcripts and digital artifacts, asynchronous votes and comments, and digital and physical artifacts from the in-person meeting. We used an inductive approach (15) to code this data: we assembled references to AI (including references to, e.g., ChatGPT) found in the artifacts and allowed thematic categories to emerge from the data.

III. RESULTS

Participants anticipate that AI, especially large language models (LLMs), will have a substantive impact on the professional practice of computing and, in turn, that this impact needs to be reflected in how CS students are instructed. One participant commented, “I think that AI [is] something that everybody in the industry is going to have to understand. [There is] the impact of the software development process. The other [impact] is you’re going to get called upon to apply machine learning to problems that you’re working on.” This comment points to two different kinds of impact that participants envision for AI: AI impacts the software development process since it is used as a *tool* by developers (e.g., GitHub Copilot), and AI is also a distinct *subdomain* of CS, similar to cybersecurity. In addition to perceiving AI as a tool and as a CS subdomain, a third theme emerged: AI is something of a *wildcard*, and it is too soon to understand the precise contours of its impact or the best ways to incorporate it into CS instruction. These three themes – AI as tool, subdomain, and wildcard – structure our discussion of this project’s findings.

A. Tool

There was broad consensus that AI tools would become prominent in CS. One focus group participant noted that

the “ability to critically collaborate with generative AI is a good skill to have.” Similarly, another participant explained that “There’s a skill in asking ChatGPT or anything how to ask it for what it is that you’re looking for. How you ask the generative AI, that’s a skill by itself.”

Due to the power of AI tools, participants saw AI as expanding access to computing: “I imagine in 20 years, everybody’s engaging in computing in some way that’s more deep than today. Even if it’s just through [a] generative AI chatbot.” One participant conceptualized these AI tools as a new way to bridge the gap between computational thinking and CS: “One of the ways in which the world changes is code does not have to be the representation of how we engage with the computer. I think a lot of the challenges with computational thinking as a field separate from computer science is you can only get so far before you have to program in order to see the fruits of your labor. It’s possible that [generative] AI makes that easier.” But these practices will lead to a need for deeper skills in some areas of computing, as another participant explained: “I feel like the . . . articulation of exactly how you want something solved is going to be even more important. . . with something like AI, it can take a well-articulated problem and a well-articulated direction of a solution and then be able to fill in the details, but it still takes someone who’s giving that orchestration of exactly how the problem is solved and thinking through and articulating really well exactly what they want.” One participant envisioned a future computing workforce where there workers who “are not strictly developers [but that] absolutely will be using things like generative AI every day” leading to a situation where most tech workers have a shallow knowledge of AI with a few having a deeper understanding, including “the pitfalls of it so that they can better pick out when things go wrong.” This statement launched more discussion of what AI expertise would look like:

First Speaker: That leads us to the question of ‘what skills do students need to understand generative AI to be that person in the next 5, 10 years on the team who actually can debug it, can understand its limitations?’ I think that circles us back to the . . . the math, stats, probability, linear algebra side of things, as well as distributed computing, distributed systems.

Second Speaker: I was imagining . . . domain expertise in whatever the generative AI is attempting to do. Which those are two different things that I think you need both of. You need somebody who can ensure that the system isn’t as inscrutable as it often feels like. Then you also still need whatever domain expertise we’re relying on the generative AI to do. You need somebody who can be critical of that and ensure that we are integrating it in effective ways.

The ability of AI tools to make computing more broadly accessible did lead to concern about how to categorize

instruction related to using AI; one focus group participant noted that “I don’t know if [the ability to use generative AI is] more computer science or is it more what I would say digital citizenship or just more of how do you best use online resources in general.”

In addition to seeing AI as a tool used to support the expansion of computing, participants also were concerned about inappropriate student use of LLMs. One of the changes anticipated by higher education CS faculty was that policies in CS courses would need to explicitly address whether and how LLMs could be used in the course. Several participants felt that some recent changes to the AP CS Principles exam were driven by fear that students would use an LLM this way; they thought that these changes were poorly designed, perhaps due to the speed with which they had to be implemented.

B. Subdomain

Participants felt that AI was quickly becoming an important subdomain of computing. Focus group participants anticipated that one of the coming changes in higher ed CS would be that students majoring in CS would be required to specialize in a CS subdomain, such as AI or cybersecurity; they also anticipated that CS concepts would be integrated into other subjects at the college level, such as training in how to use ChatGPT in English classes. As one focus group member articulated it: “it would even be better if you could integrate it with some of the other subjects.” Many participants felt that AI education should be included in foundational high school CS courses. However, when participants in the in-person meeting were asked not just to generate CS course content but to *prioritize* it, most didn’t include AI. Those who did classed it as a peripheral concern or mentioned solely its impacts and influence (not its technical aspects). A few others referred more generally to the need to teach about emerging technologies, which would include AI. There was no consensus across participant groups on the (lack of) prioritization of AI, however: one of the highest vote earners from asynchronous participants on what CS content to prioritize in a high school CS course was AI as a branch of CS.

One concern raised was that incorporating AI as a subdomain of CS would result in an unmanageable increase in the course content. One participant observed that they felt that there should be AI content, but “I just don’t know how to squeeze all that in.” Participants also related the technical skills needed to understand AI to its social impacts: “You need to know somewhat how it’s working to know how the biases slip in. . . . The understanding of that you’re scooping up data and information from a society that has biases in it already, and you just have to be aware of that.”

C. Wildcard

One participant noted that AI education was necessary but that “we’re kind of afraid of it, we’re not covering it,”

despite the fact that it already is “a big part of their life.” Nonetheless, the general sense of most participants was that “we have to embrace it and live with it because it is pretty amazing.” Participants expressed uncertainty about several aspects of AI education, reflecting the quickly-morphing nature of the technology. How to organize AI content in CS courses was one area what lacked clarity, with one participant observing that it wasn’t clear how CS courses should incorporate AI; at the in-person meeting, some participants included AI as a subset of data and analysis content; they did not create a discrete category for AI content. Others did view AI as its own topic area.

Participants in the industry focus group were asked to explore what CS knowledge, skills, and dispositions would be needed by industry in five, ten, and twenty years. At both the five and ten year time frame, there was wide agreement that the ability to use generative AI would be an in-demand skill. However, their perceptions of the landscape at the twenty-year mark were different: they did not believe that AI would be an in-demand skill. In fact, most items mentioned at longer time frames were more general and less technical skills, such as the ability to think critically and research a problem and its solutions. There was also hesitation about what CS content would be needed two decades in the future, with one participant noting that “It’s challenging to think about” and another noting that “it’s hard to imagine” because it is so difficult to predict the directions in which technology will develop. One participant noted that it was unclear to what extent generative AI would be able to replace traditional computer programming. One of the few points of conflict that arose regarding AI was the extent to which it might replace human expertise in general:

First Speaker: Industry [wants] that person to be an expert and be able to solve any problem the minute they start paying them, and it’s not based on reality. Pretty soon it will be with the use of AI because everyone will be an expert on everything without really needing to be an expert.

Second Speaker: I don’t know about that.

First Speaker: Yes, I’m certain.

Second Speaker: I’m well aware of that. [chuckles]

There was also a lack of clarity on how to mitigate problems caused by AI. For example, in the midst of discussion of biased AI training data, one focus group participant asked, “what’s the skill set to combat it? If it is bias in generative AI, then what do you need to learn to be effective in working against those biases?”

IV. DISCUSSION AND CONCLUSION

As described above, we found that participants viewed AI, in the context of CS education, as a tool, as a subdomain, and as a wildcard. There was broad consensus that new AI tools, particularly generative AI and LLMs, would have a substantive impact on CS education and

practice. We offer several recommendations for AI in CS education based on the findings of this project.

One: Address Gaps in the Discourse Some issues related to AI were *not* widely discussed may suggest topics that merit further attention. For example, there was no direct emphasis on potential problems with AI such as data privacy (16) and environmental impact (17), although these might have been implicitly included under the umbrella of AI's societal impact. While AI's impact on *learning* was explored at length, its impact on *teaching* was not. The possibility of AI to disrupt education has received less public attention than, for example, its potential in transportation or health care (18), although some groups, such as Chile's Ministry of Education, have explored how LLMs might be used to facilitate active learning (19) and there is substantive research on AI-based intelligent tutoring systems in CS instruction (20).

Two: Focus on Equity The AI toolkit mentioned previously (3) highlights the importance of addressing equity concerns related to AI tools, and several participants felt that AI instruction needs to interweave its technical and social facets. This meshes well with other research calling for an interdisciplinary approach to AI instruction (21).

Three: Promote Explainable AI One participant noted that they “suspect in future conversations, especially surrounding curriculum and AI, the statement ‘it can’t be [an opaque] box’ will be a recurring theme.” Many new AI tools do not provide any information to help the user understand *how* the tool generated its output, although methods for explainable AI do exist (22).

Four: Plan for Uncertainty As described above, focus group members found it quite difficult to project industry needs twenty years into the future. It is worth noting that, today's high school students will be only in the middle of their careers in twenty years, necessitating CS (and AI) instruction that can prepare students for a future with many unknowns. Uncertainty in the face of rapidly-developing AT technologies will also likely be an issue for curriculum design and assessment. For example, in the fall of 2023, some CS faculty struggled as Replit, a commonly used online IDE, added AI functionality that violated many plagiarism policies; Replit then abruptly ended their educational IDE (23).

Five: Seek Clarity and Consensus Participants discussed AI as a tool and as a CS subdomain, but there was some overlap and lack of clarity at the boundary between the two, particularly whether and how using an AI tool would be considered CS or digital literacy. Lack of clarity on this issue may lead to problems in the development of clear standards, course descriptions, and curriculum alignment. It might be useful to think of generative AI prompts as an (ultra) high level programming language, which would allow the community to draw on the historical trajectory of programming languages to better understand the likely trajectory of generative AI and its impacts.

REFERENCES

- [1] T. Michaeli, R. Romeike, and S. Seegerer, “What students can learn about artificial intelligence—recommendations for K-12 computing education,” in *IFIP World Conference on Computers in Education*. Springer, 2022, pp. 196–208.
- [2] F. Pedro, M. Subosa, A. Rivas, and P. Valverde, “Artificial intelligence in education: Challenges and opportunities for sustainable development,” 2019.
- [3] Code.org, CoSN, Digital Promise, European EdTech Alliance, Larimore, J., and PACE, “AI Guidance for Schools Toolkit,” <https://www.teachai.org/toolkit>, 2023, [Accessed 20-11-2023].
- [4] W. Liang, M. Yuksekgonul, Y. Mao, E. Wu, and J. Zou, “GPT detectors are biased against non-native English writers,” *Patterns*, vol. 4, no. 7, p. 100779, Jul. 2023. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S2666389923001307>
- [5] California Department of Education, “Learning with AI, Learning about AI,” <https://www.cde.ca.gov/pd/ca/cs/aiincalifornia.asp>, n.d., [Accessed 20-11-2023].
- [6] A. Klein, “Building AI literacy: What one district has learned,” <https://www.edweek.org/technology/building-ai-literacy-what-one-district-has-learned/> 2023/04, 2023, [Accessed 20-11-2023].
- [7] A. A. de Freitas and T. B. Weingart, “I’m going to learn what?! Teaching artificial intelligence to freshmen in an introductory computer science course,” in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, 2021, pp. 198–204.
- [8] J. D. Rodríguez-García, J. Moreno-León, M. Román-González, and G. Robles, “Evaluation of an online intervention to teach artificial intelligence with learningml to 10-16-year-old students,” in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, 2021, pp. 177–183.
- [9] H. Burgsteiner, M. Kandhofer, and G. Steinbauer, “Irobot: Teaching the basics of artificial intelligence in high schools,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, 2016.
- [10] M. E. Vachovsky, G. Wu, S. Chaturapruek, O. Russakovsky, R. Sommer, and L. Fei-Fei, “Toward more gender diversity in CS through an artificial intelligence summer program for high school girls,” in *Proceedings of the 47th ACM Technical Symposium on Computer Science Education*, 2016, pp. 303–308.
- [11] D. Touretzky, C. Gardner-McCune, F. Martin, and D. Seehorn, “Envisioning AI for K-12: What should every child know about AI?” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 9795–9799.
- [12] Z. Tkáčová, L. Šnajder, and J. Guniš, “Artificial intelligence—a new topic in computer science curriculum at primary and secondary schools: Challenges, opportunities, tools and approaches,” in *2020 43rd International Convention on Information, Communication and Electronic Technology (MIPRO)*. IEEE, 2020, pp. 747–749.

- [13] S. Shin, "A study on the framework design of artificial intelligence thinking for artificial intelligence education," *International Journal of Information and Education Technology*, vol. 11, no. 9, pp. 392–397, 2021.
- [14] G. K. Wong, X. Ma, P. Dillenbourg, and J. Huan, "Broadening artificial intelligence education in K-12: Where to start?" *ACM Inroads*, vol. 11, no. 1, pp. 20–29, 2020.
- [15] D. R. Thomas, "A general inductive approach for analyzing qualitative evaluation data," *American Journal of Evaluation*, vol. 27, no. 2, pp. 237–246, 2006. [Online]. Available: <https://doi.org/10.1177/1098214005283748>
- [16] CDT; Center for Democracy & Technology, "Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI," Center for Democracy and Technology, United States of America, Tech. Rep., Sep. 2023. [Online]. Available: <https://policycommons.net/artifacts/4859805/executive-summary/>
- [17] M. C. Rillig, M. Ågerstrand, M. Bi, K. A. Gould, and U. Sauerland, "Risks and benefits of large language models for the environment," *Environmental Science & Technology*, vol. 57, no. 9, pp. 3464–3466, 2023.
- [18] D. Schiff, "Out of the laboratory and into the classroom: the future of artificial intelligence in education," *AI & Society*, vol. 36, no. 1, pp. 331–348, 2021.
- [19] Ministry of Education, "Guía para docentes: Cómo usar ChatGPT para potenciar el aprendizaje activo," <https://ciudadaniadigital.mineduc.cl/wp-content/uploads/2023/05/Guia-para-Docentes-Como-usar-ChatGPT-Mineduc.pdf>, 2023, [Accessed 20-11-2023].
- [20] R. E. Francisco and F. de Oliveira Silva, "Intelligent tutoring system for computer science education and the use of artificial intelligence: A literature review." *CSEDU (1)*, pp. 338–345, 2022.
- [21] J. Henry, A. Hernalesteen, and A.-S. Collard, "Teaching artificial intelligence to K-12 through a role-playing game questioning the intelligence concept," *KI-Künstliche Intelligenz*, vol. 35, no. 2, pp. 171–179, 2021.
- [22] P. Giudici and E. Raffinetti, "Shapley-lorenz explainable artificial intelligence," *Expert systems with applications*, vol. 167, p. 114104, 2021.
- [23] T. Mann, "Compsci teachers panic as Replit kills off educational IDE," Nov 2023. [Online]. Available: https://www.theregister.com/2023/11/21/comp_sci_panic/