Jiang, B. et al. (Eds.) (2025). Proceedings of the 33rd International Conference on Computers in Education.
Asia-Pacific Society for Computers in Education

Centering Youth Perspectives in Agent Design: Piloting an Activity to Scaffold Discussions on Ethics & Agents

Namrata SRIVASTAVA[a#*], Sarah BURRISS[b#], Ashwin T.S.[a], Clayton COHN[a], Naveed MOHAMMED[a], Umesh TIMALSINA[a], Menton DEWEESE[c], Angela EEDS[c], Gautam BISWAS[a]

[a]*Department of Computer Science, Vanderbilt University, United States*
[b]*Department of Teaching and Learning, Vanderbilt University, United States*
[c]*Collaborative for STEM Education and Outreach, Vanderbilt University, United States*
*namrata.srivastava@vanderbilt.edu
#These authors contributed equally to this work

Abstract: In this paper, we share the design and preliminary analysis of student responses to a collaborative whiteboard activity centered around ethical design of a pedagogical agent (CoPA) in a STEM + computing learning environment. We report on qualitative analysis of student perceptions of (1) possible risks and benefits of using AI agents like CoPA; (2) trust & transparency; (3) consent & control; (4) fairness & bias; (5) privacy; and (6) AI learning & literacy. Though several areas of consensus across groups emerged in student responses, further investigation of relationships between agent design and function and the nature and quality of human-AI relationships is needed. We discuss (re)design of this activity, offering this pilot as a useful but imperfect structure for design discussion and seeking student input.

Keywords: AI literacy, collaborative learning, K-12 education, AI ethics, pedagogical agent, participatory design, privacy, bias, trust

1. Introduction

Students are increasingly exposed to AI-backed tools and processes in the classroom, but they rarely have opportunities to meaningfully contribute to designing or critiquing them. Co-design and participatory design processes offer a theory and method for incorporating student input into the design process, but remain under-utilized particularly in terms of *ethical* design and application of classroom technologies. This paper reports on an activity design using collaborative whiteboard software to scaffold ethics and design discussions with a group of high school students. In the following sections, we review literature on pedagogical agent design and student involvement in AI and AI ethics discussions and policy development, then describe a collaborative whiteboard activity designed to foster discussions about ethical design of AI agents while gathering information about students' perspectives in order to inform design processes. In this work, we sought to understand intersections between design and ethics, examining student responses to ethical design questions after a session interacting with an agent in a learning environment.

2. Related Work

2.1. Pedagogical Agent Research, Theory, & Design

Pedagogical agents are computational characters that can support learners by providing guidance and encouragement, demonstrating principles, and offering examples. They are designed to adapt to learners' processes and meet their evolving needs (Martha and Santoso, 2019). They have been studied for over two decades and continue to evolve along with technological and theoretical developments. Early work, such as Baylor and Kim (2005), identified and operationalized three instructional roles—Expert, Motivator, and Mentor—and demonstrated how these roles positively impact both learning and motivation. Subsequent systematic reviews confirm that pedagogical agents can positively influence learning outcomes, motivation, and affective engagement, though the magnitude of these effects remains modest and context-dependent (Martha & Santoso, 2019; Schroeder, et al., 2025). The theoretical underpinnings of pedagogical agent design draw on social agency theory, which posits that social cues—such as voice, facial expression, and embodiment—encourage learners to engage more deeply with instructional material (Mayer, 2005; Schroeder et al., 2025). Additionally, distributed cognition theory positions agents as external cognitive tools, extending learners' problem-solving capacities by scaffolding reasoning and action (Zhang & Patel, 2006). Together, these perspectives frame agents as not only communicative partners but also as distributed resources for cognition.

Design dimensions such as embodiment, modality, and persona have been shown to shape learners' perceptions and outcomes. Pedagogical agents vary widely in form—textual chatbots, voice-based assistants, 2D/3D avatars, or anthropomorphic characters. Evidence suggests that role and persona (e.g., tutor vs. peer/companion) are especially influential in determining learners' affective engagement and persistence (Baylor & Kim, 2005; Schroeder et al., 2025). For example, tutor agents often emphasize direct instruction, while co-learner agents foster collaboration and shared inquiry. Studies in K–12 contexts suggest that agents can be particularly effective at providing encouragement and scaffolding metacognitive processes, especially when aligned with learners' developmental needs (Martha & Santoso, 2019). Recent umbrella reviews highlight unresolved design questions. While text-based agents sometimes outperform narrated agents in terms of facilitating comprehension, animated or richly embodied agents do not always provide significant learning gains (Schroeder et al., 2025). This suggests that surface-level realism or animation is less important than functional alignment of the agent's role, communication style, and task context. Ongoing research therefore emphasizes tailoring design to learners' goals, contexts, and cultural settings, while foregrounding ethical considerations about trust, bias, and over-reliance on AI-supported guidance.

2.2 AI Ethics, AI Design, & Youth

Exploring the ethics and social impact of artificially intelligent technologies is often listed as a core component of "AI literacy," and the term has gained traction in both popular and scholarly discourse; researchers and practitioners in educational spaces have increasingly acknowledged the importance of engaging with ethics and social impact as a key part of learning about AI (Casal-Otero et al., 2023; Ng et al., 2021). However, youth remain largely left out of conversations about AI design, AI policy, and ethical AI use--including in their own schools. In the United States, many schools lack clear AI policies (NCES, 2024), much less systematic ways to incorporate student input. Notable exceptions to this trend include (albeit

often-isolated) efforts to include young people in policy discussions (e.g, UNICEF's AI For Children Project & Council of Europe's (2020) Declaration on Youth Participation in AI Governance), in classroom activities oriented around exploring students' AI rights (e.g., Burriss et al., 2024), and in research studies where youth (re)design AI (Chang et al., 2022; 2024) or weigh in on its ethical implications (e.g., Higgs & Stornaiuolo, 2024). AIED designers and researchers also incorporate youth perspectives to varying degrees, especially in work that uses participatory and co-design methodologies (e.g., Lim et al., 2025; Pu et al., 2025). For example, Humburg et al. (2024) describe their process of incorporating youth perspectives into narrative-centered computer learning environments. In the following sections, we share one approach to gathering youth perspectives on ethical design of an AI-backed pedagogical agent, including an activity design and student responses.

3. Method

3.1 Study and Activity Description

C2STEM and CoPA: C2STEM (Hutchins et al., 2020) is a collaborative, computational modeling environment designed to support integrated learning of physics and computing concepts. Within C2STEM, students construct, test, and debug block-based computational models to simulate scientific phenomena, thereby engaging in both disciplinary content learning and computational thinking practices. Here we focus on C2STEM's *Truck Task*, an activity targeting one-dimensional kinematics. In this task, students are asked to collaboratively model the motion of a truck that starts from rest, accelerates to 15 m/s, and then cruises at a constant velocity near 15 m/s for the remainder of the simulation until it exits the visible screen area. To complete the challenge, students must apply and integrate their knowledge of acceleration, velocity, and displacement while simultaneously leveraging computational modeling practices such as abstraction, debugging, and iterative refinement.

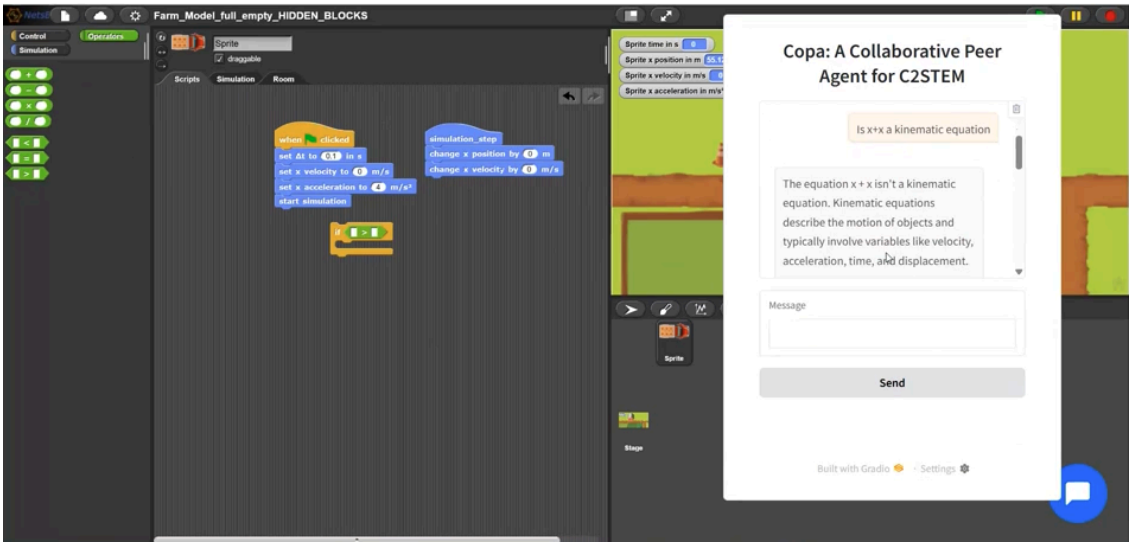

Figure 1. Snapshot of a student's screen as they interact with CoPA.

As part of the activity, students had access to *CoPA* (Collaborative Peer Agent), a GPT-4o–based conversational agent embedded within the C2STEM environment (see

Figure 1) that uses log contextualized-retrieval augmented generation (LC-RAG) (Cohn et al., 2025). CoPA was designed to function as a peer-like collaborator rather than a solution provider: it could observe students' modeling progress, prompt reflection, and offer guidance. Students were informed that CoPA might not always be correct, and that it was their responsibility to critically evaluate and decide whether to implement its suggestions. Similar to working with a human partner, students were encouraged to share sufficient details with CoPA in order to receive more contextually useful support.

Collaborative Whiteboard Activity: Following the modeling activity, students participated in a collaborative whiteboard activity using Miro, an online platform for visual collaboration. Each group worked on a shared board (see Figure 2) with a series of structured prompts that guided their reflection on the role of CoPA and, more broadly, on the use of AI agents in learning environments. The activity followed the following sequence: first, students brainstormed risks and benefits of using AI agents like CoPA. Next, students responded to guiding questions within each category, including Trust & Transparency, Consent & Control, Fairness & Bias, Privacy, and AI Learning & Literacy[1]. One group followed the board's written instruction to move any risks and benefits to their respective categories, but the others did not, instead generating all new responses in notes under each category heading. Lastly, a researcher facilitated a whole group discussion to reflect on the most important themes and discussion points raised during the whiteboard activity.

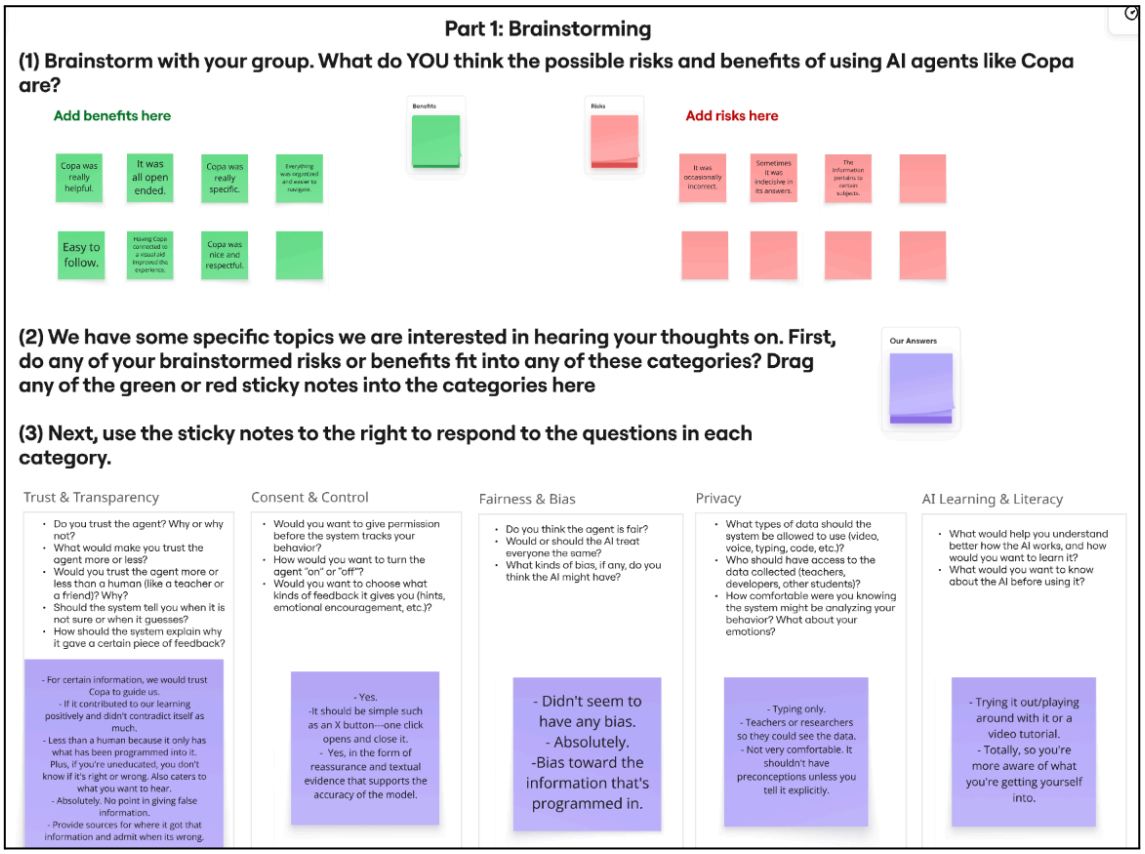

Figure 2. One student group's collaborative whiteboard space.

3.2 Participants, Data & Analysis

[1] The detailed questions used for the Miro board activity can be found here.

We conducted the study as a one-day, 2.5 hours classroom session in our lab environment with a total of 16 high school freshmen (ages 14-15). For the first task (Truck task), students worked in pairs, and later re-formed into five groups for the collaborative whiteboard activity (three groups with 4 students and two groups with 2 students). The study was approved by Vanderbilt University's Institution Review Board, and both parental consent and student assent were obtained prior to participation. We collected data from multiple sources during the study: student interaction logs from the C2STEM environment, and student-generated responses on the Miro board. For the purpose of this paper, we focus on the Miro board data, which consisted of 123 responses generated by the five groups.

For this analysis, we performed several iterative rounds of qualitative coding (Saldaña, 2021). We began with procedural coding, tagging each student response according to the specific whiteboard question it addressed. Next, we engaged in open coding, generating initial coding in a fast and exploratory manner and not confining them to any particular categories. We then conducted axial coding, grouping, refining and relating these codes to identify higher-level categories. Finally, we synthesized our findings within and across each category.

4. Findings

4.1 Student Perspectives on Possible Risks and Benefits

Student groups generated a total of 21 risks and 23 benefits related to using AI agents like CoPA. In our qualitative coding, we found several themes in their responses as we grouped related responses together. For benefits, these included many responses praising Copa for being easy to use and (potentially) helpful, with agent affordances including having access to correct answers and student code. However, students also listed accuracy as a potential risk, indicating that although student experiences with CoPA's efficacy and reliability were mixed, students could envision a way that an AI system like CoPA could be a good source of information and assistance.

Many benefits that students listed were related to how humans (both teachers and students) might work together with AI to reap learning or professional benefits. Students at two tables wrote about how an agent might help teachers (though one of these same groups also wrote "loss of jobs of educators" as a potential risk), including "makes teachers jobs easier."'Table three emphasized working with AI as "collaboration" and that it "Helps you discover solutions that you wouldn't without AI," showing how they saw AI as a way to augment human ability. Though some students listed over-reliance on AI as a potential risk, many listed benefits of the CoPA agent in terms of its open-endedness and hint-giving (rather than answer-giving) behavior, like "It guided without giving answers." This indicates to us the potential for mitigating some forms of over-reliance with these kinds of "guiding" rather than more prescriptive agent behaviors, which students responded positively to in this study.

For risks, we found that the largest thematic category of responses was related, broadly, to poor answer quality. Four out of five table groups indicated some form of concern with the quality of the responses the agent might give, variously citing potential for (or experience with the CoPA agent) giving wrong, confusing, vague, repetitive, or inconsistent answers. Aside from this common trend, we coded four other themes or singly occurring response types: cheating, over-reliance, creativity. Two tables wrote about (presumably

students) cheating on assignments as a risk, but one of these same tables also wrote "could plagiarize others' work," indicating that they saw risk of both students and AI engaging in ethically dubious authorship practices. Two tables indicated that over-reliance, "codependence," or reduction in critical thinking was a potential risk, and three mentioned loss of "creativity" or "originality" as a risk. Finally, one group included teacher job loss as a potential risk, and one mentioned "lack of understanding" as compared to humans.

4.2 Student Perspectives on Trust & Transparency

With the exception of a student who indicated they'd trust the agent "*more than a stranger less than a friend,*" groups wrote that they would trust a human (like a peer or teacher) more than the agent. One table noted that an agent "*Also caters to what you want to hear*" as a reason for trusting the agent less than a human.

Though one group responded that "*For certain information, we would trust Copa to guide us,*" others expressed distrust of the agent. One member of Group 3 wrote, *"I do not trust the agent because It has given me confusing, redundant, or incorrect answers."* In addition to distrust stemming from poor perceptions of agent performance in the activity, students also mentioned "*fallacies among AI*" and breach of privacy ("*could store my vocal info without my permission*") as reasons for their distrust. All responding tables replied *"yes"* to whether they'd want to know if the agent was guessing, with one student writing, "*that would make it 10x better*."

In describing ways that would increase their trust of the agent, students included responses like "*Provide sources for where it got that information and admit when it's wrong.*" Three groups indicated that accurate and reliable performance of the agent would be linked to increased trust, describing performance in terms of "*Correct and helpful answers*" that *"would incentivize me to trust the agent more, vice versa would make me trust it less," "responding to feedback well,"* and *"[not] contradict[ing] itself so much."*

4.3 Student Perspectives on Consent & Control

All tables did want to give permission before system tracking, and all tables also wanted to be able to turn the agent on and off, via switch, button, or clicks. Group 5 offered, *"It should be simple such as an X button—one click opens and closes it."*

Four of the table groups also offered ideas for if/how they would want to choose the kind of feedback they received from the agent, addressing both content and tone. One group wrote, *"I would want to choose the intent of the feedback and how it is given: encouraging, instructional, steps, ect [sic]."* Another echoed this idea about combining supportive ("encouraging" or "reassuarance") agent functions with procedural or evidence-giving functions: *"Yes, in the form of reassurance and textual evidence that supports the accuracy of the model."* Regarding the kind of information an agent should give, a third group offered that the agent should give *"small hints, not revealing the entire code."* The final responding group said they'd like to tailor the kind of feedback on their mood: *"Yes, depends on mood, like if you're mad you don't want fake encouragement."* However, these groups also indicated in other sections of the board that they did not want the system to read their emotions in any sort of automated way, indicating that they would want to actively choose how the agent responds based on their own (human) self-assessment of their mood. Since we worded the question around user choice, this is unsurprising, but could be a point of clarification for future studies.

4.4 Student Perspectives on Fairness & Bias

All tables felt the agent was fair. Responses were more mixed when discussing potential bias. One group had brainstormed a potential benefit related to bias that they dragged and dropped into this category (they were the one group to respond to this instruction on the collaborative board), writing that *"It can give unbiased feedback."* While one group felt the agent *"Didn't seem to have any bias,"* others thought that the agent may have *"Bias based on sources that the bot was trained on"* or *"toward the information that's programmed in."* One student mentioned that the agent may *"be ignorant to other opinions on non-factual topics."*

Only one group offered an answer that wasn't a form of *"yes"* to whether an agent should "treat everyone the same." They answered that (presumably a student interacting with the agent) *"Should take a quick pre-use quiz to see proficiency."* Others wrote, *"The AI should treat everyone the same with the goal of giving correct information"* and *"The AI must treat all users with the same respect, even if they are abusing the bot."*

4.5 Student Perspectives on Privacy

While one group wrote that access to students' information should be determined by the individual user, others thought that teachers (3 groups), researchers (2), and "*maybe other students*" (1) should be able to access the data collected by the system. One group wrote, *"I think teachers having access to the data can be beneficial to some levels, but complete control or knowledge over the data would be unneeded."*

All responses were negative toward incorporating emotion detection into the environment/agent, and two tables only wanted "*typing*" or "*typing and code, not voice or video*" data to be used by the system. One group wrote, "*A version with emotion and facial recognition would be a bit creepy"* and another implied doubt in the system's ability to respond appropriately to complex human emotion: *"Human emotions are complicated and are hard even for humans to create a response to."*

4.6 Student Perspectives on AI Learning & Literacy

3 groups responded that hands-on lessons/exploration or a tutorial would be a useful way to learn about AI and how it works. All groups wanted to know more about the system, with one providing this reasoning: *"Totally, so you're more aware of what you're getting yourself into."* The other 4 groups gave specific ideas about what they might want to know about the system before using it, including those related to **optimizing use** (i.e., "*how you can word you messages to best convey what you want*"); **agent function & development**, and what risks may follow ("h*ow they process what you input into the system*," "*What information.. the bot [is] trained on," "How much information was fed to the bot, and is there [sic] bias*," and "How the AI processes things, and where it might go wrong"); and **privacy and transparency** ("*how it tracks you*," for the system to "*Disclose what data it's collecting and why,*" "*What information is specifically being collected and where it's going*").

5. Discussion & Significance

5.1 (Re)designing Agents: Ethical Design & AI Literacy

Though this was a small sample of students, we saw consistently how central agent **design and performance** were to students' perceptions of the agent, including its ethical risks and benefits. While there were some areas of consensus (e.g., that students wanted to have control over some aspects of the agent and knowledge about how the system was built and functioned; discomfort with automated emotion detection and response; desire for the agent to indicate when "guessing"), there were also many ways that groups contradicted each other and even themselves when assessing both CoPA's functionality and ethical issues; for example, the same group that listed "*accuracy 10/10*" as a benefit of using AI agents like CoPA also wrote "*I don't trust its answers*." This kind of response merits further investigation of students' ideas about trust in relation to agent accuracy; is accuracy necessary, but insufficient, for trust? Was there disagreement within the group that was not captured in the collaborative whiteboard? Another group wrote that "*Correct and helpful answers would incentivize me to trust the agent more, vice versa would make me trust it less*." Others made similar points, pointing to providing sources and not "contradict[ing] itself as much" as ways to increase their trust in the agent. We see students connecting accuracy, helpfulness, and transparency (i.e., providing sources, indicating uncertainty) to trust, but these relationships must be explored further in order to gain insight into not only what makes an agent more trustworthy, but whether that should be the goal in the first place. For example, if an agent is almost always "*correct and helpful*," *should* a student trust it? How might trust relate to other potential risks mentioned by students, like over-reliance on AI? Especially given burgeoning concern over human relationships with chatbots, from AI-generated characters (e.g., Roose, 2024) to children's toys (Brandon, 2025), questions about our interactions with agents–and how we should design them ethically–have only become more urgent.

It is not enough to design agents to be more accurate; we must also consider the ethical dimensions of how performance is linked to children's relationships with AI. For example, an ethical design question may include, "(How much) do we want students to build trust with this agent, and how might we achieve that through model performance and critical education about the agent?" This design question is, of course, also an AI literacy question. As we create agents that engender trust (or frustration, as the case may be), we must think about carefully teaching that an agent is not human–but flawed and capable in different ways than human socialization has trained us to attend to. Addressing potential misconceptions like *"it can give unbiased feedback*" will be crucial to students' understanding not only of the agent and its role in their learning, but also of larger sociotechnical concepts like bias.

5.2 Limitations, Affordances, & Future Directions

Though the activity design is not the focus of this paper, there are some clear affordances and drawbacks to the way we designed and conducted the activity we describe here. One major limitation is our small sample size for this pilot, alongside extremely limited time students had to complete the activity. Given our time constraints, we developed this collaborative whiteboard activity with the intention of sparking small-group discussion and gathering group perspectives efficiently. However, this also meant that we missed out on individual students' perspectives, and group responses were subject to group dynamics, where a single student may have dominated response-writing. In the future, we hope to have more time to conduct ethical design discussions and also seek individual feedback via survey and/or interview with students. Additionally, this paper reports only on students'

written feedback on the collaborative whiteboards. In future work, we will incorporate analysis of student audio recordings and environment logs. Finally, while we wanted to spark conversation within the groups by offering several different questions for each topic, this made the written whiteboard results more difficult to interpret. Researchers mapped each response, where possible, to the applicable sub-question, but this was time-consuming and imperfect.

Despite these limitations, we were able to see distinct trends and disagreements among students in their responses to our ethical design questions. As activity facilitators, we also saw how students engaged with enthusiasm and interest in this discussion, sharing their perspectives with each other and with the whole group. Furthermore, attaching this kind of a discussion and reflection activity to hands-on experimentation with an agent in a learning environment served several important functions: first, it provided a concrete scaffold for having conversations about abstract ethical concepts, meaning that all students were able to participate in critical AI literacy learning; second, it provided us with important design feedback about the agent; finally, it signaled to the students that their feedback was, in fact, important to us.

6. Conclusion

The question, *How do we design agents in support of student learning that are helpful, encouraging, and accurate?* seems worthy, but it is incomplete without consideration of what is also *ethical*. In this paper, we offer that one way to move beyond design that prioritizes accuracy and efficacy to broach deep ethical issues like the nature and quality of human-AI relationships, is to start discussions about these things with the students themselves. This kind of activity has distinct affordances and drawbacks for gaining insight into students' views of agent design and ethics. While it gave us a rich glimpse into groups' ideas about these topics, we see opportunities for furthering this analysis with audio and log data; for capturing individual perspectives via survey or interview; and for redesigning this activity to apply to more agent design contexts.

Acknowledgements

We would like to thank the students who participated in this study. This work was supported by the NSF Award #DRL-2112635 and IES Award #R305C240010. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF or IES.

References

Baylor, A. L., & Kim, Y. (2005). Simulating instructional roles through pedagogical agents. *International Journal of Artificial intelligence in education, 15*(2), 95-115.

Brandon, M. (n.d.). *Will AI-Powered Toys Rewire Childhood?* | *Psychology Today*. Retrieved August 25, 2025, from https://www.psychologytoday.com/us/blog/the-future-of-intimacy/202507/will-ai-powered-toys-rewire-childhood

Burriss, S. K., Hutchins, N., Conley, Z., Deweese, M. M., Doe, Y. J., Eeds, A., Villanueva, A., Ziegler, H., & Oliver, K. (2024). Redesigning an AI bill of rights with/for young people: Principles for

exploring AI ethics with middle and high school students. *Computers and Education: Artificial Intelligence*, *7*, 100317. https://doi.org/10.1016/j.caeai.2024.100317

Casal-Otero, L., Catala, A., Fernández-Morante, C., Taboada, M., Cebreiro, B., & Barro, S. (2023). AI literacy in K-12: A systematic literature review. *International Journal of STEM Education*, *10*(1), 29. https://doi.org/10.1186/s40594-023-00418-7

Chang, M. A., Philip, T. M., Cortez, A., McKoy, A., Sumner, T., & Penuel, W. R. (2022). *Engaging youth in envisioning artificial intelligence in classrooms: Lessons learned* (Rapid Community Reports).

Chang, M. A., Wong, R. Y., Breideband, T., Philip, T. M., McKoy, A., Cortez, A., & D'Mello, S. K. (2024). Co-design partners as transformative learners: Imagining ideal technology for schools by centering speculative relationships. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–15. https://doi.org/10.1145/3613904.3642559

Cohn, C., Rayala, S., Snyder, C., Fonteles, J., Jain, S., Mohammed, N., Timalsina, U., Burriss, S. K., S, A. T., Srivastava, N., Deweese, M., Eeds, A., & Biswas, G. (2025). *Personalizing Student-Agent Interactions Using Log-Contextualized Retrieval Augmented Generation (RAG)* (No. arXiv:2505.17238). arXiv. https://doi.org/10.48550/arXiv.2505.17238

Council of Europe. (2020). *Declaration on Youth Participation in AI Governance*.

Higgs, J. M., & Stornaiuolo, A. (2024). Being Human in the Age of Generative AI: Young People's Ethical Concerns about Writing and Living with Machines. *Reading Research Quarterly*, *59*(4), 632–650. https://doi.org/10.1002/rrq.552

Humburg, M., Dragnić-Cindrić, D., Hmelo-Silver, C. E., Glazewski, K., Lester, J. C., & Danish, J. A. (2024). Integrating Youth Perspectives into the Design of AI-Supported Collaborative Learning Environments. *Education Sciences*, *14*(11), 1197. https://doi.org/10.3390/educsci14111197

Hutchins, N. M., Biswas, G., Maróti, M., Lédeczi, Á., Grover, S., Wolf, R., Blair, K. P., Chin, D., Conlin, L., Basu, S., & McElhaney, K. (2020). C2STEM: A System for Synergistic Learning of Physics and Computational Thinking. *Journal of Science Education and Technology*, *29*(1), 83–100. https://doi.org/10.1007/s10956-019-09804-9

Lim, J. O., Barkhuff, G., Awuah, J., Clyde, S., Sogani, R., Gardner-McCune, C., Touretzky, D., & Uchidiuno, J. O. (2025). Escape or D13: Understanding Youth Perspectives of AI through Educational Game Co-design. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–17. https://doi.org/10.1145/3706598.3714037

Martha, A. S. D., & Santoso, H. B. (2019). The design and impact of the pedagogical agent: A systematic literature review. *Journal of Educators Online, 16*(1), n1.

Mayer, R. E. (2005). Cognitive Theory of Multimedia Learning. In R. Mayer (Ed.), *The Cambridge Handbook of Multimedia Learning* (pp. 31–48). chapter, Cambridge: Cambridge University Press.

National Center for Education Statistics (NCES). (2024). *School Pulse Panel: Technology*. https://nces.ed.gov/surveys/spp/results.asp#technology-dec24-chart-4

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, *2*, 100041. https://doi.org/10.1016/j.caeai.2021.100041

Pu, I., Ravi, P., Dinh, L. D., Joe, C., Ogoe, C., Li, Z., Breazeal, C., & Ostrowski, A. K. (2025). "How can we learn and use AI at the same time?": Participatory Design of GenAI with High School Students. In *Proceedings of the 24th Interaction Design and Children* (pp. 204–220). Association for Computing Machinery. https://doi.org/10.1145/3713043.3727057

Roose, K. (2024, October 23). Character.ai Faces Lawsuit After Teen's Suicide. *The New York Times*. https://www.nytimes.com/2024/10/23/technology/characterai-lawsuit-teen-suicide.html

Saldaña, J. (2021). *The coding manual for qualitative researchers* (4th ed.). SAGE Publications.

Schroeder, N. L., Davis, R. O., & Yang, E. (2025). Designing and learning with pedagogical agents: An umbrella review. *Journal of Educational Computing Research, 62*(8), 1907-1936.

Zhang, J., & Patel, V. L. (2006). Distributed cognition, representation, and affordance. *Pragmatics & Cognition, 14*(2), 333–341. https://doi.org/10.1075/pc.14.2.12zha