Embedding Ethical Awareness in Computer Science and AI Education: The PEaRCE Approach to Responsible Computing

Mohammad Hadi Nezhad[1*], Francisco Enrique Vicente Castro[2], Eugene Mak[1], Peter J. Haas[1], Danielle Allessio[1], Leon Osterweil[1], Injila Rasul[1], Heather Conboy[1], and Ivon Arroyo[1]

[1] University of Massachusetts Amherst, Amherst, MA 01003, USA
{mhadinezhad}@umass.edu
[2] New York University, New York, NY 10012, USA

Abstract. Despite the widespread use of computer systems, their societal impacts are often poorly understood, highlighting the need for the AIED community to acknowledge and contribute to incorporating Ethical and Responsible Computing (ERC) into Computer Science (CS) education. We introduce the Platform for Ethical and Responsible Computing Education (PEaRCE), an interactive tool created to integrate ERC into post-secondary education through realistic workplace simulations. PEaRCE scenarios guide students through ERC dilemmas—awareness, decision-making, feedback—during the processes of developing advanced AIED and other technologies, where software/hardware may have beneficial and harmful impacts. Moreover, we integrate PEaRCE into CS courses via a sequence of structured learning modules and trained "ERC Teaching Assistants (TAs)" to support the integration process. We present preliminary insights from our deployment experience, suggesting PEaRCE's potential to enhance ERC awareness and reasoning. We discuss the possibility of embedding PEaRCE into AI in Education courses.

Keywords: Ethics in AIED · Ethical and Responsible Computing Education · Educational Technologies · Computer Science Education

1 Introduction

Advanced computer systems have become embedded into daily life, transforming how people work, learn, and interact with societal institutions and each other. While the rapid integration of computer technologies into society offers new opportunities, their societal impacts often remain poorly understood. This raises critical questions about what consequences may arise from the use and integration of these technologies into our existing societal structures or how they might amplify already-existing societal concerns.

The use of AI-driven systems in education has highlighted these challenges. For example, England and Ireland implemented algorithmic scoring systems to standardize results and prevent grade inflation [23]. However, these systems were criticized for disproportionately downgrading students from certain groups (e.g., those attending disadvantaged schools) by overemphasizing past institutional performance, causing psychological distress, perceptions of unfairness by students and families, and legal challenges and widespread protests affecting schools and authorities. Facial recognition technology, often used for security, automated attendance, and emotion detection in schools, is often argued by experts to create surveillance environments that may undermine students' autonomy and freedom [2]. These systems have also exhibited biases related to race and gender, raising concerns about unfair treatment of certain student groups [22, 28].

The AIED community acknowledges the risks of neglecting ethical considerations during the development, research, and application of AI in education, which ultimately aims for a more equitable and inclusive educational environment for the general population [20, 26]. Beyond technical soundness, designing societally responsible AI technologies should not be an afterthought but a core value instilled among researchers, engineers, and developers of current and future AI systems [18, 20]. However, computing education has historically often overlooked this aspect, and researchers have noted that engineers and developers often lack systematic education on the broader societal impacts of the technologies they create [18]. In response, scholars have explored various approaches to integrating Ethical and Responsible Computing (ERC) education into computing curricula [4, 5, 13, 21].

One approach to ERC education is to offer standalone courses that focus on ERC issues. Fiesler et al. [12], however, point out that this approach, by itself, fails to teach students that ERC considerations and technical practices are very closely intertwined. We assert, therefore, that it is crucial to embed ERC education into CS courses in a manner that deeply engages students, scales to large numbers of students and classes, and can overcome potential instructor resistance. Our approach to realizing this vision uses an educational tool—the *Platform for Ethical and Responsible Computing Education* (PEaRCE)—that immerses students in role-playing simulations.

PEaRCE is a web-based interactive simulation tool designed to embed ERC education into CS courses at post-secondary educational institutions. PEaRCE engages students in realistic workplace scenarios where they apply their technical knowledge (e.g., AI, data science, software development) in complex, ambiguous situations involving the development of systems that may produce tangible personal or societal benefits but could potentially cause physical, economic, psychological, or societal harm. Students are presented with the opportunity to explore technology's implications via conversations with simulated stakeholders, helping them increase their awareness of ERC issues they may face in the workplace and learn to seek out diverse and informative stakeholders for responsible decision-making. Reflection questions throughout the scenario deepen their engagement with the scenario as it unfolds.

To further increase engagement and effectiveness, we provide supplementary curricular materials in the form of pre- and post-simulation activities. Following prior work [12], we use trained graduate students (“ERC-TAs”) to coordinate with instructors on deploying PEaRCE and its accompanying activities.

In this paper, we describe PEaRCE and examples of its simulation scenarios (Section 3), our protocol for integrating PEaRCE into CS classes (Section 4), and our experiences deploying PEaRCE over the past year in undergraduate CS courses at a large US public university, along with preliminary insights from pre- and post-test surveys and interviews on PEaRCE’s impact on students ERC awareness and reasoning (Section 5). We conclude with a discussion of future work in the context of AI in Education. Our work thus far suggests that the PEaRCE approach has the potential to train future AIED developers and researchers on ERC education. This has critical implications for *AI education* and *AI in education* as the AIED community continues to think about what it means to design AI systems and how to prepare students to contribute to the design of socially-beneficial technologies.

2 Background and Related Work

There is a growing consensus among AI researchers that developing AI systems should be grounded on robust and ethical AI principles and universal human rights and obligations [8]. Scholars have attempted to define factors that are essential when considering designing AI for social good, including—beneficence (promoting human and planetary well-being); non-maleficence (avoiding harm, misuse, and privacy violations); autonomy (preserving human freedom in decision-making); justice (ensuring fairness, preventing discrimination, and fostering diversity); and explicability (ensuring transparency and accountability) [14].

The ethics of AI in Education is likewise gaining increasing attention, as seen in the increasing number of articles on Ethics and AIED within the AIED community in recent years [16]. This represents a shift from an assumption of the field’s self-perception as inherently ‘good’ because of its educational purposes [26] and starts addressing the need for awareness and oversight in AIED technologies upholding the above principles.

Holmes and colleagues [20] claim that most AIED researchers lack the training to tackle emerging ethical challenges in AI for education. They suggested that various stakeholders—including developers, educators, and policymakers—to be provided with key information about the pros and cons of specific AIED technologies, including known limitations and benefits that are likely to emerge from the use of specific AIED systems. However, for AIED developers to recognize these impacts, they must first acquire knowledge and skills about ethical awareness and reasoning. *Who* trains future AIED researchers and developers to develop such mindsets and practices thus becomes a critical point of discussion.

Several efforts have been made to integrate ethics into post-secondary computing curricula, with a common strategy of creating stand-alone ethics courses for teaching ethics and AI. This approach has shown its prevalence in university-level technology and ethics course syllabi [17]. However, scholars also argue that separating ethics into its own courses disconnects it from the actual practice of computing; instead, ethical considerations should be integrated into technical courses to recognize their intertwined connection [12].

A review of 100 papers on ethics education in computing courses by Brown et al. [4] found that 62 supported integrating ethics into existing courses. This integration ranges from a single ethics module in technical courses to multiple ethics-focused modules. The authors suggested that the level of integration should align with learning goals, the intended conception of ethics, and context. For instance, while stand-alone courses may be better for teaching formal ethical frameworks and principles, embedding ethics in technical courses helps students identify ethical issues within specific technical implementations.

While directly embedding ethics in computing courses is promising, doing this effectively remains a challenge. One major obstacle is that many instructors lack formal ethics training and/or knowledge and may feel unqualified to teach it [4, 12, 24, 27]. In addition, already overloaded curricula require an extensive workload to incorporate ethics. Computing educators also need expertise in pedagogical and instructional design to effectively embed ethics interventions in their coursework [4, 7]—including designing suitable readings/lectures, in-class discussions, assignments and projects, and ethics awareness and reasoning assessment methods to measure learning outcomes and evaluate the effectiveness of specific interventions.

Educators have also incorporated various ethics-related content such as ethical principles [25] and ethical frameworks (e.g., deontology, consequentialism) [9, 27]. It is also argued that simply using an ethical framework or ethical principles does not guarantee ethical decision-making and that computing education should also consider ways individuals make ethical decisions (i.e., behavioral ethics) [19]. The author therefore advocates fostering students' critical thinking, exposing them to scenarios where computing professionals' decisions can indirectly harm individuals, building students' confidence to speak up and make their voice and values heard—especially in unsupportive organizational environments or under external pressures (see also [18]), and to help them feel comfortable navigating ethical ambiguities and to consider different aspects of a situation. The latter can be approached by engaging students with diverse technical and non-technical stakeholders to consider different viewpoints [11, 21, 25, 27]. Researchers also stress using real-world examples—e.g., news articles and personal stories—to make ethical issues tangible. Feffer et al. [11] noted that students often struggle to envision AI system failures in real-world scenarios without prior exposure, advocating for educational tools that expose students to past incidents long before they start their career as AI practitioners.

Student engagement with ethical topics remains another challenge. Researchers have explored various strategies such as interactive simulations [3, 6], game-based learning [15], and encouraging faculty-led discussions on responsible computing, which can help students engage more seriously with ethical content [7]. Real-world case studies can further enhance engagement by connecting ethical issues to practical and tangible applications [11, 12]. Some prior work also argued that strong technical foundations are a prerequisite for meaningful engagement with ethical topics, especially in complex computing courses [10, 12]. Overall, prior studies suggest combining these strategies to optimize student engagement.

3 Platform for Ethical and Responsible Computing Education (PEaRCE)

Developed at a large US university, PEaRCE is an interactive online platform that simulates workplace scenarios involving ERC issues, where students navigate the simulation by making choices that shape the unfolding scenario. Each scenario begins by introducing a context, typically a proposed product or project whose stated goal is usually corporate profit, market positioning, or societal benefit (Fig. 1, top-left). Ethical or societal drawbacks are subtly hinted at but not emphasized. The scenario introduces the student's role, typically through a job description (Fig. 1, top-middle), and asks the student whether to accept the position or explore its ramifications by engaging in conversations with stakeholders potentially impacted by the proposed technology—such as bosses, colleagues, experts, lawyers, friends, and family members (Fig. 1, bottom-left). The simulation includes reflective prompts, encouraging students to comment on their understanding of the situation, their conversational choices, and the ethical or technical issues they face. These reflections enhance the learning experience and provide instructors with insights to gauge students' learning.

The goal is for students to sense the scenario's societal downsides and gain confirmation by engaging with stakeholders who have the most to say about those issues. Students can converse with a limited number of stakeholders (e.g., 5 out of 14), encouraging thoughtful selection of the most valuable interactions (Fig. 1, bottom-middle). A conversation is considered effective if it touches key ethical concerns, involves a directly affected stakeholder, and/or provides insights from a domain expert. During these interactive conversations, students choose questions from a list and see how the stakeholder responds (Fig. 1, bottom-right).

After the conversations, students are again given the choice of proceeding, or delaying the project for further discussion with stakeholders. Upon choosing to go ahead with the project, they receive feedback including a summary of the ethical quandary, hypothetical outcomes of proceeding, modifying, or abandoning the project, and a discussion of one's responsibility to ask critical questions and anticipate potential positive and negative consequences. Some hypothetical outcomes are intentionally unpleasant—such as public backlash over invasive data collection or biased interpretations of a foreign accent—to emphasize the seriousness of their decisions and the importance of careful consideration.

Students receive additional feedback in the form of a radar plot, showing how thoroughly their chosen conversations have touched upon the various issues and highlighting which issues are more ethically germane (e.g., user privacy concerns over promotion opportunities). The chart shows how students prioritized stakeholders within the conversation limit and what they missed by not selecting certain discussions. Finally, students are provided with URL links to real-world examples of similar ethical quandaries and resources such as the ACM Code of Ethics and Professional Conduct.

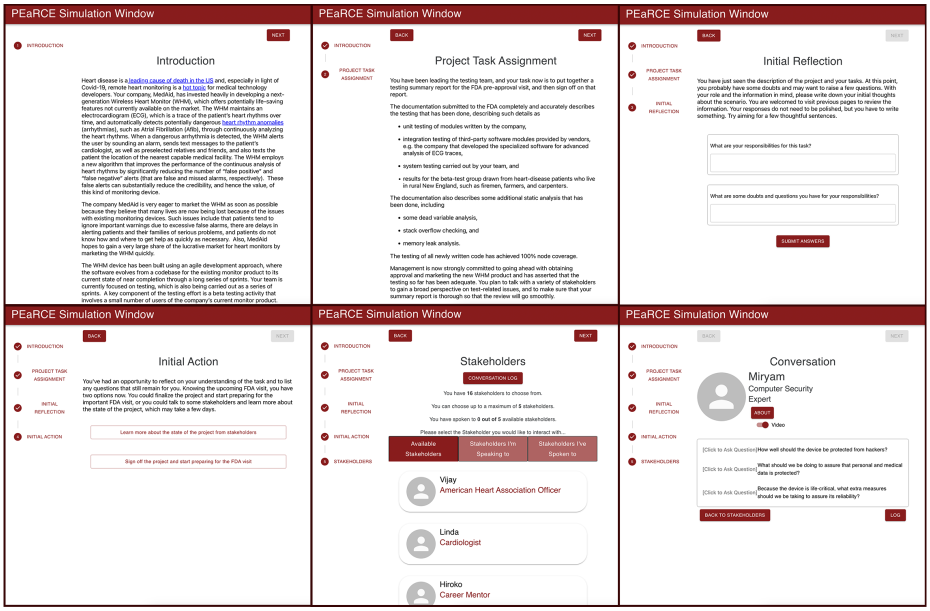

Fig. 1. Sample PEaRCE Interface Stages: (top-left) introduction to the context/technology, (top-middle) project task assignment (student role), (top-right) initial reflection on the project and tasks, (bottom-left) initial decision point, (bottom-middle) list of stakeholders, and (bottom-right) sample conversation with a stakeholder.

3.1 Examples of PEaRCE Work Scenarios

PEaRCE is a scenario-based simulation designed for scalability, allowing future expansions and updates to meet the needs and contexts of different courses across universities. The PEaRCE team is collaborating with volunteer CS students to develop work scenarios that align with course content, making scenarios and ethical discussions more relevant and engaging for both students and teaching staff. Below, we present three example scenarios.

Scenario 1: Heart Monitoring Device Testability. The context of this scenario is the imminent arrival of a review team from the US Food and Drug Administration (FDA) at a company seeking FDA approval of an innovative heart monitoring device. Among the innovations is the ability of the device to not only detect heart rhythm anomalies, but then to send emails about the event to family or caregivers, to identify nearby doctors and emergency facilities and automatically send the help deemed appropriate. The student is placed as the leader of the testing team who is to make a presentation intended to convince the FDA that the device has been adequately tested. The student believes that the testing has been quite comprehensive but has qualms about the thoroughness of the testing of some of the communications capabilities. Stakeholder conversants include bosses, colleagues, testing experts, doctors, communications experts, college friends, and parents. A persuasive presentation could lead to potential improvements in patient care, as well as raises, promotions, and favorable visibility in a growing market. Downsides include violation of privacy, embarrassing and wasteful emergency calls, marketplace failure (and possible lawsuits) due to product defects that might have been detected by more thorough testing.

Scenario 2: My Medical Advisor. The context of the My Medical Advisor scenario is a company, currently operating a medical information/advice platform, that has used data gathered from customer website interactions—typing speeds, typing errors, repetitive requests, etc.—to develop a machine learning model that detects cognitive decline in older adults. The company wants to sell lists of customers flagged by the model to vendors, who can provide appropriate health products and services, and to the Department of Motor Vehicles, who can remove dangerous drivers from the road. The job description presented to the student is to lead the effort to identify and gather additional data for improving the model prior to the sales effort. Stakeholder conversants include bosses and colleagues, but also lawyers, privacy advocates, parents, and data mining experts. Doing this job well could lead to improved patient outreach and care and safer roads, as well as raises, promotions, and prestige. Downsides include violation of privacy and model biases leading to incorrect lists that could result in harassment and inappropriate revocation of drivers licenses.

Scenario 3: VR in K-12 STEM Education. This scenario is about a new marketing push by the maker of a Virtual Reality (VR) headset into the education market. The company is repurposing the VR headset and adapting it for use in schools, perhaps even in elementary schools. The student's job is to develop software to make the VR headset particularly attractive to the students in these schools by creating a recommender system aimed at bringing appropriate materials to the right students, measuring their responses to these materials, and then adjusting the recommender system to increase student and teacher satisfaction. Stakeholders include company bosses and executives, work colleagues, child education experts, recommender system experts, child psychologists, college friends, and parents. Success with this job should lead to exciting educational opportunities for students, as well as raises, promotions, admiration by colleagues, and greater job mobility. Downsides include the inadvertent inclusion of inappropriate materials in recommended lists, exacerbation of extant problems caused by excessive screen time, dizziness, and difficulty in distinguishing reality.

4 Integrating PEaRCE into Computer Science Courses

Integrating PEaRCE into a class as a standalone assignment would dilute its impact and create a perceived disconnect between ERC education and technical material. Instead, we supplemented the simulation with multiple activities throughout the semester, fostering deeper student engagement, peer discussions, and reflection. Seamless integration of these activities into technical CS classes was a key challenge. Following Cohen et al. [7], we trained a cohort of Ethical and Responsible Computing Teaching Assistants (ERC-TAs). An ERC-TA is a graduate student who supplements regular course TAs and collaborates with both the PEaRCE team and course staff to manage the integration of PEaRCE into the curriculum, as described later in this section.

One key advantage of using trained ERC-TAs is improved instructor acceptance. Prior research has highlighted that instructors' lack of training and discomfort in teaching ERC, along with reluctance to add ERC content to an already crammed syllabi, can lead to instructor resistance [4, 12, 24, 27]. By delegating some responsibilities to ERC-TAs and limiting in-class activities to one session, our approach reduces instructor workload while addressing concerns about class time loss. Interestingly, as described in Section 5.2, some students stated a desire for more in-class activities. The specific sequence of PEaRCE activities within a course is outlined below.

Pre-test Survey: Assessing Students' Initial Awareness. For research purposes, we collaborated with course staff and designed a pre-test survey to evaluate students' initial ERC awareness, ensuring that the survey aligns with the course schedule, workload, and timing. The survey is designed to be integrated as an early take-home course assignment and involves reading three newspaper articles (200-400 words each), adapted from the case studies of the ACM Code of Ethics and Professional Conduct [1]. Each article provides context for a publicly used technology (e.g., medical implant, internet content filtering), covering aspects such as the business project, technology features, and related incidents—all fictional and designed for educational purposes. While maintaining a balanced tone, the articles differ in emphasis: one elaborates more on the negative aspects, another on positive aspects, and the third remains neutral. Each article is followed by two open-ended questions. The first question prompts students to identify the potential parties or communities positively impacted by the technology and describe the associated possible benefits (i.e., pros). The second question follows the same format, focusing instead on the possible negative impacts (i.e., cons). The survey also asks students for feedback on the clarity of the articles and the two follow-up questions to refine future iterations.

Video Lecture: Priming Students for PEaRCE. A few weeks after the pre-test survey, students watch a recorded lecture by a PEaRCE faculty member. It emphasizes that past harmful systems have not infrequently been developed by well-meaning people, and some of those people may well have been young people, perhaps recent graduates from CS programs. The intent here is to make students aware that societally worrisome systems may arise from systems built by people like them. These points are driven home by presenting examples, such as the Volkswagen diesel engine system, whose software had been designed to deploy its pollution-suppression system only when the vehicle was being tested. While senior executives may have planned the deception, lower-level developers wrote and tested the code. The video stresses developers' responsibility to critically assess their work's impact and voice their concerns. The video ends by introducing the PEaRCE project and its goals. Overall, the lecture aims to spark curiosity around ERC and prepare students for upcoming PEaRCE activities.

Using the PEaRCE Platform. Around the mid-semester, an ERC-TA introduces students to the PEaRCE platform via a very short in-class presentation (around 5-10 minutes), providing an overview of its functionality to ensure they are prepared to engage with the simulation. The ERC-TA then registers students in the system (via auto-generated invitation emails) and monitors their activity for grading purposes. While students engage with the web-based PEaRCE simulation as a take-home assignment, they can report any questions or technical issues to the teaching staff through the course's communication channels.

In-Class Discussions: Reflecting on the Experience. In this phase, an ERC-TA leads a 45-75 minute class session, guiding students through a reflective discussion on their experiences with the PEaRCE scenario. To support ERC-TAs in facilitating a productive session, the PEaRCE team provides a template slide deck that incorporates aggregated student log data from the PEaRCE platform and lays out key discussion topics, including (1) ERC issues students encountered during the scenario, (2) summary statistics on student decisions, e.g., the percentage of students who decided to engage with stakeholders versus proceeding with their assigned task, (3) the number of students who engaged with each of the stakeholders, with an accompanying discussion around whether, in hindsight, the students would modify their choices, (4) the simulation's potential impact on students' future professional development and decision-making practices, and (5) students' real-life experiences where they could have evaluated the impacts of their decisions more thoroughly. At the end of the session, students receive additional resources to support their learning, such as the ACM Code of Ethics and Professional Conduct, ethics-related university courses, and opportunities to engage with student groups focused on ethics and technology.

Post-test Survey: Assessing Learning Gains. The post-test survey aims to evaluate how PEaRCE involvement influences students' awareness of ERC. Later in the semester, students complete the same assignment as the pre-test survey, allowing for a direct comparison of their responses.

5 PEaRCE Deployment Experience

This section describes our pilot deployment from Spring 2024, detailing the courses involved and the scenarios used (Section 5.1). We then discuss preliminary insights from the deployment (Section 5.2). Our research protocols, including research consent, study procedures, data collection, and analysis, were approved by our institutional IRB.

5.1 Courses and Scenarios

We deployed PEaRCE in two different CS courses offered in the Spring 2024 semester. One course was a 300-level undergraduate Introduction to Software Engineering course. PEaRCE was deployed in two sections, each taught by a different instructor. Participation in PEaRCE activities contributed 2-3% of the total course credit, and one ERC-TA was assigned to each section. We used the Heart Monitoring Device Testability simulation scenario (Section 3.1) to align with the course content.

The second course was a master-level industry mentorship course offered in a single section. PEaRCE activities accounted for 5% of the total course credit, and one ERC-TA was assigned. Since the course focuses on real-world data science practices, such as data collection and use of machine-learning-based analytics, we selected the My Medical Advisor simulation scenario (Section 3.1) to align with the course content.

5.2 Exploration and Preliminary Insights

Insights from Student Responses to Pre- and Post-test Surveys. We conducted a preliminary thematic analysis of 24 students' responses to pre- and post-test surveys (Section 4) from one section of the Software Engineering course where PEaRCE was deployed. Our analysis focused on student responses regarding the potential benefits (pros) and harms (cons) of the technology solutions presented in the newspaper articles. A primary coder identified pros and cons within the dataset, and a second coder re-coded 30% of responses to ensure high inter-rater reliability (inter-class correlation, $R = 0.98$, $p < 0.001$). A third researcher tallied the total number of potential harms and benefits in each response, and results are reported based on these averages.

Figure 2 shows an increase in both the variety and quantity of potential benefits and harms that students identified. Notably, they significantly increased their recognition of the technologies' potential harms (paired samples t-test, $t(24) = -3.05$, $p < 0.006$). Overall, students' responses focused on two main areas: (1) technological aspects—e.g., privacy, security, data handling, bugs, and (2) societal impacts including health/wellness, economic/financial, legal, trust, and accessibility.

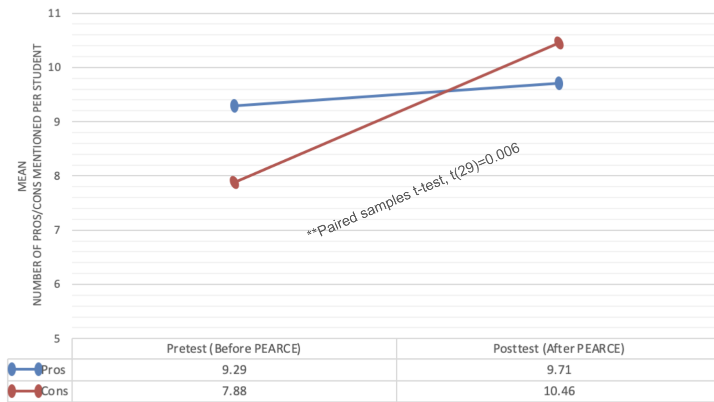

Fig. 2. Mean Pros and Cons reported by students before and after PEaRCE

Insights from Interviews. We conducted semi-structured interviews in early summer after the Spring 2024 semester. Participants included five students from three course sections, two of the three ERC-TAs, and one course instructor. Each interview lasted 30-60 minutes and was led by a primary interviewer and a facilitator, both part of the PEaRCE research team. All interviews were conducted remotely via Zoom, recorded (video and audio), and transcribed for analysis.

Our preliminary insights suggests that students and ERC-TAs viewed stakeholder conversations in the PEaRCE simulation positively. Students found them engaging: *"I was pretty engaged in [conversations] and interested in learning about what their position was on these situations"* and that they bring a broader perspective to the course content: *"[...] we usually only learn about coding it efficiently, but we don't really look at how it affects the customer [...]"*.

We also received several suggestions for making PEaRCE activities more engaging. Both ERC-TAs and students suggested more in-person involvements, with one student noting: *"in-person interaction versus a recording might've been more helpful so people could ask questions in real-time"*. While the simulation-based implementation was interesting, long texts were sometimes less engaging: *"some of the explanations [...] was a bit long and time-consuming to read every example that was given"*. An ERC-TA suggested more instructor involvement: *"[...] if it's the professor that becomes [involved], they [students] treat it a little more seriously"*. Another key insight for our multi-modular deployment was that engagement in one activity can highly depend on previous activities: *"[some students] got a lot out of PEaRCE or at least did it seriously and found it interesting, and wanted to talk about it [during in-class discussions] versus maybe [those] who didn't have much to say [...]"*.

Feedback on PEaRCE's connection to course content was mixed. Students appreciated the video lecture covering broader topics beyond CS: *"[the lecturer] did not restrict himself to just CS [...] he gave examples from different walks of life and gave us a really good idea of what to expect in real-world scenarios"*. However, they suggested deeper integration with the course: *"if it could be com-*

bined more deeply with the course, I think that would definitely be better". Some also recommended inviting guest speakers: *"[...] to bring in either researchers or students who have more experience introducing the ethical concepts [...]"*. One student suggested allowing scenario selection based on relevance to their work: *"[...] different scenarios and I could probably choose based on which is more relevant to me"*. ERC-TAs valued real-world connections, with one noting: *"the more relevant the content is to the real world, the more effective it can get"*, and another describing the differences between simulation and real-world complexity: *"when you're making that decision and you know that this is an ethics simulation, you kind of know what the right choice is in that sense. But in the real world, I think it's a lot more nuance, right?"*

6 Discussion, Conclusion & Future Work

We argued that AIED systems should not be assumed to be ethically designed or inherently beneficial simply because they serve an educational purpose. Instead, the AIED community should actively promote ethical and responsible computing (ERC) education for both current and future practitioners, including students, researchers, engineers, and developers. This education should take place well before these individuals begin their careers—e.g., during post-secondary studies—so they develop the practical knowledge, skills, and confidence needed to consider and voice diverse perspectives in technology development, especially when ethical and social concerns are overlooked or inadequately considered.

With this in mind, we introduced PEaRCE (Platform for Ethical and Responsible Computing Education), an interactive tool designed to promote ERC education through simulations of realistic workplace scenarios where advanced technologies are developed. In the simulations, users navigate complex situations, voice their opinions at reflection points, make critical decisions—e.g., engaging with stakeholders or proceeding with assigned jobs—, and get feedback on the potential consequences of their decisions. We also outlined our multi-modular approach to integrating PEaRCE into CS courses, drawing from pedagogical research, instructional design, and the expertise of our interdisciplinary team.

We deployed PEaRCE in CS courses at a large US public university. Preliminary insights suggests PEaRCE's potential to enhance students' ability to recognize and articulate ethical concerns when analyzing case studies on computing. Interviews revealed that students responded positively to PEaRCE and actively engaged with its different modules. They also provided suggestions to further improve the experience, such as incorporating more in-person interactions, multimedia content, and further integrating PEaRCE and course content.

Our position is that PEaRCE has the potential to train future AIED professionals on ethical aspects of developing educational systems—e.g., through our sample scenario on virtual reality headset design. We aim to expand PEaRCE by refining its current scenarios and developing additional ones around the ethical challenges of emerging AIED technologies, particularly those that may directly or indirectly impact key educational stakeholders (e.g., teachers, students, parents).

By doing so, we seek to equip future AIED professionals with the knowledge and skills to make well-informed decisions toward ethical and responsible computing.

As we prepare to expand PEaRCE beyond our university, we invite the AIED community to collaborate with us in refining the platform, enhancing existing scenarios, creating new ones, and adapting PEaRCE for other institutions by considering their educational structures and values. The long-term goal of the PEaRCE team is bold: We aim to make a profound impact by advancing Ethical and Responsible Computing training for the next generation of computer technologists, benefiting society as a whole.

Publisher's Acknowledgement This version of the contribution has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-031-98414-3_10. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use https://www.springernature.com/gp/open-science/policies/accepted-manuscript-terms

References

1. ACM: Case studies: ACM integrity project (nd), https://ethics.acm.org/integrity-project/case-studies/
2. Andrejevic, M., Selwyn, N.: Facial recognition technology in schools: critical questions and concerns. Learning, Media and Technology **45**(2), 115–128 (2020). https://doi.org/10.1080/17439884.2020.1686014
3. Barab, S., Dede, C.: Games and immersive participatory simulations for science education: An emerging type of curricula. Journal of Science Education and Technology **16**(1), 1–3 (2007)
4. Brown, N., Xie, B., Sarder, E., Fiesler, C., Wiese, E.S.: Teaching ethics in computing: A systematic literature review of ACM computer science education publications. ACM Transactions on Computing Education **24**(1) (2024). https://doi.org/10.1145/3634685
5. Castro, F., Raipura, S., Conboy, H., Haas, P., Osterweil, L., Arroyo, I.: Piloting an interactive ethics and responsible computing learning environment in undergraduate CS courses. In: Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1. p. 659–665. SIGCSE 2023, ACM (2023). https://doi.org/10.1145/3545945.3569753
6. Chernikova, O., Heitzmann, N., Stadler, M., Holzberger, D., Seidel, T., Fischer, F.: Simulation-based learning in higher education: A meta-analysis. Review of Educational Research **90**(4), 499–541 (2020). https://doi.org/10.3102/0034654320933544
7. Cohen, L., Precel, H., Triedman, H., Fisler, K.: A new model for weaving responsible computing into courses across the CS curriculum. In: Proceedings of the 52nd ACM Technical Symposium on Computer Science Education. p. 858–864. SIGCSE '21, ACM (2021). https://doi.org/10.1145/3408877.3432456
8. Corrêa, N.K., Galvão, C., Santos, J.W., Del Pino, C., Pinto, E.P., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., de Oliveira, N.: Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. Patterns **4**(10), 100857 (2023). https://doi.org/10.1016/j.patter.2023.100857

9. Cowan, B., Fourquet, E., Kryven, M.: Teaching the societal consequences of computer science: new ideas for increasing student involvement. In: Proceedings of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education. p. 242–247. ITiCSE 2018, ACM (2018). https://doi.org/10.1145/3197091.3197146
10. Cowan, B., Fourquet, E., Kryven, M.: Teaching the societal consequences of computer science: new ideas for increasing student involvement. In: Proceedings of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education. p. 242–247. ITiCSE 2018, ACM (2018). https://doi.org/10.1145/3197091.3197146
11. Feffer, M., Martelaro, N., Heidari, H.: The AI incident database as an educational tool to raise awareness of AI harms: A classroom exploration of efficacy, limitations, & future improvements. In: Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization. EAAMO '23, ACM (2023). https://doi.org/10.1145/3617694.3623223
12. Fiesler, C., Friske, M., Garrett, N., Muzny, F., Smith, J.J., Zietz, J.: Integrating ethics into introductory programming classes. In: Proceedings of the 52nd ACM Technical Symposium on Computer Science Education. p. 1027–1033. SIGCSE '21, ACM (2021). https://doi.org/10.1145/3408877.3432510
13. Fiesler, C., Garrett, N., Beard, N.: What do we teach when we teach tech ethics? A syllabi analysis. In: Proceedings of the 51st ACM Technical Symposium on Computer Science Education. p. 289–295. SIGCSE '20, ACM (2020). https://doi.org/10.1145/3328778.3366825
14. Floridi, L., Cowls, J.: A unified framework of five principles for AI in society. Machine Learning and the City: Applications in Architecture and Urban Design pp. 535–545 (2022)
15. de Freitas, A.A., Weingart, T.B.: 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. p. 198–204. SIGCSE '21, ACM (2021). https://doi.org/10.1145/3408877.3432530
16. Fu, Y., Weng, Z.: Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered ai practices. Computers and Education: Artificial Intelligence **7**, 100306 (2024). https://doi.org/https://doi.org/10.1016/j.caeai.2024.100306
17. Garrett, N., Beard, N., Fiesler, C.: More than "if time allows": The role of ethics in AI education. In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society. p. 272–278. AIES '20, ACM (2020). https://doi.org/10.1145/3375627.3375868
18. Hagendorff, T.: The ethics of AI ethics: An evaluation of guidelines. Minds and Machines **30**(1), 99–120 (2020). https://doi.org/10.1007/s11023-020-09517-8
19. Hedayati-Mehdiabadi, A.: How do computer science students make decisions in ethical situations? Implications for teaching computing ethics based on a grounded theory study. ACM Transactions on Computing Education **22**(3) (Jun 2022). https://doi.org/10.1145/3483841
20. Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S.B., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.I., et al.: Ethics of AI in education: Towards a community-wide framework. International Journal of Artificial Intelligence in Education pp. 1–23 (2022)
21. Horton, D., McIlraith, S.A., Wang, N., Majedi, M., McClure, E., Wald, B.: Embedding ethics in computer science courses: Does it work? In: Proceedings of the

53rd ACM Technical Symposium on Computer Science Education - Volume 1. p. 481–487. SIGCSE 2022, ACM (2022). https://doi.org/10.1145/3478431.3499407
22. Karan, B., Angadi, G.R.: Potential risks of artificial intelligence integration into school education: A systematic review. Bulletin of Science, Technology & Society **43**(3-4), 67–85 (2023). https://doi.org/10.1177/02704676231224705
23. Kelly, A.: A tale of two algorithms: The appeal and repeal of calculated grades systems in England and Ireland in 2020. British Educational Research Journal **47**(3), 725–741 (2021). https://doi.org/https://doi.org/10.1002/berj.3705
24. Martin, C.D.: The case for integrating ethical and social impact into the computer science curriculum. In: The Supplemental Proceedings of the Conference on Integrating Technology into Computer Science Education: Working Group Reports and Supplemental Proceedings. p. 114–120. ITiCSE-WGR '97, ACM (1997). https://doi.org/10.1145/266057.266131
25. Petelka, J., Finn, M., Roesner, F., Shilton, K.: Principles matter: Integrating an ethics intervention into a computer security course. In: Proceedings of the 53rd ACM Technical Symposium on Computer Science Education - Volume 1. p. 474–480. SIGCSE 2022, ACM (2022). https://doi.org/10.1145/3478431.3499275
26. Porayska-Pomsta, K., Holmes, W., Nemorin, S.: The ethics of AI in education. In: Handbook of Artificial Intelligence in Education, pp. 571–604. Edward Elgar Publishing (2023). https://doi.org/10.4337/9781800375413.00038
27. Saltz, J., Skirpan, M., Fiesler, C., Gorelick, M., Yeh, T., Heckman, R., Dewar, N., Beard, N.: Integrating ethics within machine learning courses. ACM Transactions on Computing Education **19**(4) (Aug 2019). https://doi.org/10.1145/3341164
28. Stark, L.: Facial recognition is the plutonium of AI. XRDS **25**(3), 50–55 (2019). https://doi.org/10.1145/3313129