Understanding and Engaging Critical Resistance to AI in Education

Dinesh Ayyappan*
Social and Responsible Computing Research Group
Universitat Pompeu Fabra
Barcelona, Catalunya, Spain
dinesh.ayyappan@upf.edu

Davinia Hernández-Leo
DTIC
Universitat Pompeu Fabra
Barcelona, Barcelona, Spain
davinia.hernandez-leo@upf.edu

Robert Larson
Singapore American School
Singapore, Singapore
rlarson@sas.edu.sg

Leaf Elhai
Learnlife
Barcelona, Spain
leaf@learnlife.com

Julia Chatain
Dept. of Humanities, Social and Political Sciences
ETH
Zurich, Switzerland
Faculty of Mathematics and Computer Science
UniDistance Suisse
Brig, Switzerland
jchatain@ethz.ch

Grace Li
Computer Science
University of Chicago
Chicago, Illinois, USA
graceli458@uchicago.edu

Abstract

The push to integrate Generative AI technologies (GenAI) into education is framed by its advocates as an inevitable leap towards efficiency and personalization, but adoption has been met with resistance from learners and educators on the ground, who express deep concerns about deskilling and the erosion of trust and intellectual agency. This resistance is not mere technophobia or self-preservation; it is a legitimate response to a product that has landed in educational institutions worldwide. We reframe this friction not as an obstacle but as a design resource. We will bring together HCI researchers, designers, and educators to explore deliberate non-use, productive friction, or technological refusal as valid design goals. To find GenAI's place in education, we must understand both where it belongs and where it doesn't. We will engage critical resistance to AI as a starting point towards a grounded research agenda for a more human-centered future in education.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; *Interaction design*; • **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → **Interactive learning environments**.

Keywords

artificial intelligence, generative AI, K-12 education, higher education, critical thinking, critical analysis, human-centered, workshop

ACM Reference Format:
Dinesh Ayyappan, Davinia Hernández-Leo, Robert Larson, Leaf Elhai, Julia Chatain, and Grace Li. 2026. Understanding and Engaging Critical Resistance to AI in Education. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3772363.3778761

*Corresponding Author

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/26/04
https://doi.org/10.1145/3772363.3778761

1 Why this workshop?

The push to integrate recent artificial intelligence (AI) technologies into education is framed by advocates as an inevitable leap towards efficiency and personalization [13], but adoption of these products has met with resistance from students and teachers who express deep concerns about deskilling and the erosion of trust and intellectual agency [26]. This resistance is not mere technophobia or self-preservation; it is a legitimate response to a family of products that has landed in educational institutions worldwide. These products are difficult to capture with terminology [13] but for the purposes of this workshop include popular chatbots (like ChatGPT, Gemini, and Claude), productivity tools integrated into Microsoft and Google Suites, and pedagogical products like Khanmigo and LearnLM. These commercially available generative AI products and others have the potential to tackle longstanding challenges in education [12] like Bloom's 2-Sigma problem [4, 24] and the dream of one-to-one tutoring, and they also move billions of dollars in subscriptions and enterprise software [21]. This is a high-stakes battleground that must be addressed and acted on cooperatively.

This workshop takes the view that Generative AI (GenAI) is a powerful technology that likely has benefits for education [2], but that the interests of educators and students need to be more strongly

reflected in the way GenAI products are used in schools. Failure to involve stakeholders in the design process can result in irrelevant or inefficient solutions [1], so we welcome critical input as a design resource to help us chart the best way forward. If we can get this right, we can do more than just reject hype. We can make progress towards persistent challenges in education (e.g. in accessibility and inclusion [7] or learning in open-ended problems [18]). By leaning into intentionality, we aim to reduce incidental harms and elevate what might otherwise have been 'happy accidents' to design goals.

1.1 Generative AI's effects on teaching and learning

Since the release of OpenAI's ChatGPT, students' use of GenAI has become more widespread. In a 2024 global survey, 86% of students are regularly using generative AI in their studies [9], and more than 2 out of 3 students are using AI for information search and summarizations [9, 22]. But the utility of using these GenAI products was often mediated by hallucinations or inaccuracies [17, 23]. These inaccuracies affect the quality of students' work and may directly impact their grades, credibility, and learning [15]. In a large-scale field experiment in a high school, students allowed to use GenAI performed better on math practice questions but then performed worse once the AI was removed, suggesting that students had bypassed the effortful cognitive processes required for learning when practicing with GenAI products [3]. In student-AI learning scenarios, a study by Darvishi et al. found that students had a tendency to rely on rather than learn from AI [10]. These findings demonstrate that while many students have adopted GenAI technologies into their academic workflows, students often rely on these technologies as a "crutch" instead of a learning aid [3, 10].

For educators, current research is exploring how GenAI tools might support various aspects of teaching, from designing curriculum to grading student essays and even their own professional development [16]. One study found that GenAI, with the expert guidance of a teacher, can help scaffold middle school mathematics curricula by making the math problems accessible to students working below grade level [20]. On the other hand, when it came to GenAI products providing feedback on student writing, their feedback often lacked higher order feedback beyond sentence level edits [19]. In terms of content that educators teach, some argue that educators should focus on teaching the skills that GenAI cannot perform like experiential knowledge, ethics, and interpersonal relationships [6], illustrating how generative AI is not only changing how educators teach, but may change what they teach as well. These questions of how and when educators should use GenAI to augment their abilities is closely entangled with concerns that these technologies might replace them [16].

1.2 Critical resistance to Generative AI in Education

The current landscape of GenAI in education calls for more than just keeping a ledger of benefits and harms. There is a field of critical scholarship on educational technology that reminds us that the recent boom is really just the latest chapter in a long history of AI hyping [25].

The term "AI" has its origins in the 1950s as a marketing phrase to sell research and attract funding [13, 21]. Decades of progress may make resistance seem futile, and some may reject the possibility of charting a path independent of agendas set by industry [14]. While technological ventures often gain momentum and grow more difficult to challenge over time [5], there is historical precedent for the critical rejection of technology.

Author Ted Chiang reminds us that the Luddites were not anti-technology for the sake of being anti-technology but fought against factory owners for economic justice [8]. More recent is the story of Amsterdam and its relationship to cars [11]: Car usage spiked after World War II, and emissions and road fatalities began to emerge as public concerns. In the 1990s, the country passed national policies to simultaneously restrict the infrastructure allocated to cars and expand the infrastructure allocated to bicycles and public transport. Public information and education campaigns followed, and today, bicycles and cars coexist governed by policies that hold public safety at the center. GenAI in education may be akin to the Dutch car of the mid-to-late 20th century. It is time to organize a response to GenAI in education that prioritizes the learner, but to do so, we need to better understand the current situation and the nature of the policy and infrastructure levers we have at hand.

1.3 Motivating questions

We gather researchers and educators motivated to discuss questions like these over two sessions described below (subsection 4.2):

- In academic settings, where should the ideal boundaries be between where GenAI technologies are welcome and where they are not? How do we bridge the gap between this ideal and the present situation?
- What are the most relevant pedagogical considerations guiding the use of GenAI technologies? What can we transfer from how we currently teach about tools and information sources, and what is distinctly different?
- What metrics and methodologies are most useful for assessing the effectiveness of GenAI technologies in learning contexts?
- How can we adjust our design processes, or develop novel design processes, to ensure that the GenAI-powered tools we design for education are relevant, meaningful, and inclusive?
- The enterprise of education has many stakeholders. How do we balance the major voices (e.g. students, parents, teachers, administrators, industry) in institution-level decisions and policies?

2 Call for Participation

We invite researchers, educators, and designers to discuss, from a place of self-determination, the role of Generative AI technologies (GenAI) in education. Some advocates welcome a seemingly inevitable leap towards efficiency and personalization, but they have been met with resistance from students and educators on the ground. We frame this opposition as a design resource, an invaluable one, as we explore deliberate non-use, productive friction, or technological refusal as valid design goals. To find GenAI's place in education, we ask these questions:

- In academic settings, where should the ideal boundaries be between where GenAI products are welcome and where they are not?
- What are the most relevant pedagogical considerations guiding the use of GenAI products?
- What can we transfer from how we currently teach about tools and information sources, and what is distinctly different?
- What metrics and methodologies are most useful for assessing the effectiveness of GenAI products in learning contexts?
- How can we adjust our design processes, or develop novel ones, to ensure that the GenAI tools we design for education are relevant, meaningful, and inclusive?
- How do we balance the major stakeholder voices (e.g. students, parents, teachers, administrators, industry) in institution-level decisions and policies?

We ask participants to submit either a **position paper** or **research paper** addressing any of these motivating questions. These contribution pathways align with the goals of each of the workshop's two, discussion-oriented sessions, and a subset of accepted papers will be selected for lightning talks. Submissions are limited to a maximum of 4 pages (excluding references) in ACM Primary Article format (link), and can be submitted through the workshop website (link).

Participants are also asked to submit a short personal statement of their relevant background (up to 150 words) to help organizers curate a diverse set of perspectives. One or two authors of each accepted submission must attend the workshop.

3 Workshop Structure

3.1 Workshop length

This workshop will span two 90-minute sessions to allow for the meaningful conversations needed to build a constructivist understanding of the topic and subsequent research agenda.

3.2 Organizers

This workshop is organized by both secondary school educators and HCI researchers who, together, have extensive experience with the practical and theoretical issues around AI in Education.

Dinesh Ayyappan (he/him) is a PhD student at Universitat Pompeu Fabra (Barcelona), where he is researching Human-AI interaction from the lenses of safety and fairness. Prior to this, he was a high school teacher for 10 years and holds Master's degrees in Computer Science and in Secondary Education. He has won multiple awards for his teaching and has coached and ran workshops for adults on topics ranging from competency-based grading to building culturally inclusive environments.

Grace Li (she/they) is a PhD student at the University of Chicago. Her research explores the impacts that AI literacy interventions have on student-AI interactions, specifically whether teaching students about responsible AI use actually translates to students utilizing these techniques when completing a task. She is also researching the drivers and barriers educators face when integrating AI literacy into their own classrooms.

Leaf Elhai (she/her) is a literacy educator who has taught for 12 years across four countries. She currently serves as a learning guide at Learnlife, an innovative school in Barcelona where learners have the agency to explore their passions and design their own learning journey. Leaf is curious about AI as a tool for equity in the classroom but wants to ensure that AI does not replace learners' critical thinking and creativity.

Rob Larson (he/him) is an educator with 17 years of international teaching experience and a Masters degree in Technology in Learning Design. Currently a research and humanities teacher at Singapore American School, he has a front-row seat to how teachers, parents, administrators, and students are thinking about and using AI every day.

Julia Chatain (she/her) is a senior researcher at the Professorship for Learning Sciences and Higher Education at ETH Zürich, Switzerland, and an upcoming Assistant Professor in Learning Technologies and Distance Education at UniDistance, Switzerland. Her interdisciplinary research bridges Human-Computer Interaction and Learning Sciences, focusing on the design of inclusive technological solutions to improve STEM education. Specifically, she is interested in how embodied, multi-modal, and concrete learning experiences can address both cognitive and emotional challenges faced by learners.

Davinia Hernández-Leo (she/her) is Full Professor in the Department of Information and Communications Technologies at Universitat Pompeu Fabra (Barcelona) and holds Serra Húnter / ICREA Academia fellowships. She leads the TIDE research group (Interactive and Distributed Technologies for Education) and her work centers on human-centered learning technologies, learning design, CSCL, learning analytics, and AI in Education.

4 Activities

4.1 Before the Workshop

As part of the selection process, organizers will align submissions under the motivating questions of the workshop (subsection 1.3). These contributions will be summarized into a 2-page brief that will be used to ground and focus each session.

Separately, accepted submissions will be posted on the public workshop website, and a Discord server will be made for transparent communication with the organizers. A few accepted submissions that are deemed highly relevant to the session questions will be selected for short lightning talks during each session.

4.2 During the Workshop

This is a bridge-building workshop in that we want to connect educators to researchers and transition from the present towards a more ideal future. For that to work, we need to quickly build trusting, constructive relationships, so this workshop will maximize time for facilitated, focused small-group discussions over two sessions.

After periods of structured discussion, groups will share out to the larger gathering. Whether this is done visually (e.g. on posters), orally, or digitally (e.g. on Discord) is a decision for the facilitators to make closer to the workshop.

4.2.1 Session 1: Where do AI products belong in schools and universities? After a warm-up activity to get a sense of who is in the room (e.g. sharing Likert-scale opinions on debatable statements),

participants will gather in small groups with facilitators spread throughout or circulating the room. These groups will process the reading materials (subsection 4.1) and individuals will share reactions and questions as a starting point. We are prepared to use structured protocols to help jumpstart discussion.

After surfacing strengths and concerns about AI tools in education, we will pause for lightning talks about critical adoption of AI in learning environments, and then take a short break. Participants will then brainstorm educational processes and products where AI is either particularly well- or poorly-suited. Facilitators will help to prioritize a few of these that are more fraught or ought to be explored more deeply in our second session.

4.2.2 Session 2: How do we study, design, and reshape interactions to move closer to our ideal (from Session 1)? This session will begin with lightning talks focused on innovative interactions. Then, we will explore the prioritized areas from Session 1 to establish pressing research questions, experimental designs, and compelling positions to carry our thinking beyond the workshop. Participants will self-select their group based on the prioritized areas from Session 1 to increase engagement and connections that can last beyond the workshop.

Groups will contribute to a common document that will be used as the basis for a workshop report meant to be shared with the wider HCI in Education community. We will end with a total participation activity that asks everyone to share how their thinking has changed over the workshop.

4.3 After the Workshop

An informal workshop report will be compiled shortly after the workshop to maintain momentum and encourage continued discussions. If there is interest among participants, this report will be expanded to a formal publication.

5 Accessibility

5.1 Publishing plans

Submitted papers will be linked on the workshop website, and a workshop report (subsection 4.3) will be shared shortly after the workshop.

5.2 Offline materials

Interested community members who missed the session will be able to access the papers and reading materials on a workshop website. Additionally, a workshop report will be prepared that could evolve into a position paper, journal article, or special issue call.

References

[1] Riordan Alfredo, Vanessa Echeverria, Yueqiao Jin, Lixiang Yan, Zachari Swiecki, Dragan Gašević, and Roberto Martinez-Maldonado. 2024. Human-centred learning analytics and AI in education: A systematic literature review. *Computers and Education: Artificial Intelligence* 6 (2024), 100215.

[2] Shatha N Alkhasawneh, Georgios Lampropoulos, and Davinia Hernández-Leo. 2025. Generative AI in learning design: A systematic review. In *Lecture Notes in Computer Science*. Springer Nature Switzerland, Cham, 17–30.

[3] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman. 2025. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences* 122, 26 (2025), e2422633122.

[4] Benjamin S Bloom. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational researcher* 13, 6 (1984), 4–16.

[5] Michel Callon. 1990. Techno-economic networks and irreversibility. *Sociol. Rev.* 38, 1_suppl (May 1990), 132–161.

[6] Lydia Cao and Chris Dede. 2023. Navigating a world of generative AI: Suggestions for educators. *The next level lab at harvard graduate school of education* 5, 2 (2023).

[7] Julia Chatain, Gerd Kortemeyer, and Andreas R Fender. 2024. AIt: Intention-preserving Automatic Alt-text Generation for Educational Content. In *European Association for Research on Learning and Instruction (EARLI) SIG 6 & SIG 7 Conference*.

[8] Ted Chiang. 2023. Will A.I. Become the New McKinsey? *New Yorker* (May 2023).

[9] Digital Education Council. 2024. What students want: Key results from DEC Global AI Student Survey 2024. https://www.digitaleducationcouncil.com/post/what-students-want-key-results-from-dec-global-ai-student-survey-2024

[10] Ali Darvishi, Hassan Khosravi, Shazia Sadiq, Dragan Gašević, and George Siemens. 2024. Impact of AI assistance on student agency. *Computers & Education* 210 (2024), 104967.

[11] Gladys Frame, Arturo Ardila-Gomez, and Yang Chen. 2017. The kingdom of the bicycle: what Wuhan can learn from Amsterdam. *Transp. Res. Procedia* 25 (2017), 5040–5058.

[12] Michail Giannakos, Roger Azevedo, Peter Brusilovsky, Mutlu Cukurova, Yannis Dimitriadis, Davinia Hernandez-Leo, Sanna Järvelä, Manolis Mavrikis, and Bart Rienties. 2024. The promise and challenges of generative AI in education. *Behav. Inf. Technol.* (Sept. 2024), 1–27.

[13] Olivia Guest, Marcela Suarez, Barbara Müller, Edwin van Meerkerk, Arnoud Oude Groote Beverborg, Ronald de Haan, Andrea Reyes Elizondo, Mark Blokpoel, Natalia Scharfenberg, Annelies Kleinherenbrink, Ileana Camerino, Marieke Woensdregt, Dagmar Monett, Jed Brown, Lucy Avraamidou, Juliette Alenda-Demoutiez, Felienne Hermans, and Iris van Rooij. 2025. Against the Uncritical Adoption of 'AI' Technologies in Academia. doi:10.5281/zenodo.17065099

[14] Maarten Hajer and Jeroen Oomen. 2025. *Captured Futures: Rethinking the Drama of Environmental Politics*. doi:10.1093/9780198955382.001.0001

[15] Emily L Howell and Dominique Brossard. 2021. (Mis) informed about what? What it means to be a science-literate citizen in a digital world. *Proceedings of the National Academy of Sciences* 118, 15 (2021), e1912436117.

[16] Regina Kaplan-Rakowski, Kimberly Grotewold, Peggy Hartwick, and Kevin Papin. 2023. Generative AI and teachers' perspectives on its implementation in education. *Journal of Interactive Learning Research* 34, 2 (2023), 313–338.

[17] Yoonsu Kim, Jueon Lee, Seoyoung Kim, Jaehyuk Park, and Juho Kim. 2024. Understanding users' dissatisfaction with chatgpt responses: Types, resolving tactics, and the effect of knowledge level. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 385–404.

[18] John Kos, Dinesh Ayyappan, and Ashok Goel. 2024. A Constructivist Framing of Wheel Spinning: Identifying Unproductive Behaviors with Sequence Analysis. In *Generative Intelligence and Intelligent Tutoring Systems*, Angelo Sifaleras and Fuhua Lin (Eds.). Springer Nature Switzerland, Cham, 174–187.

[19] Christopher Mah, Mei Tan, Lena Phalen, Alexa Sparks, and Dorottya Demszky. 2025. From Sentence-Corrections to Deeper Dialogue: Qualitative Insights from LLM and Teacher Feedback on Student Writing. EdWorkingPaper No. 25-1193. *Annenberg Institute for School Reform at Brown University* (2025).

[20] Rizwaan Malik, Dorna Abdi, Rose Wang, and Dorottya Demszky. 2025. Scaffolding middle school mathematics curricula with large language models. *British Journal of Educational Technology* 56, 3 (2025), 999–1027.

[21] Alva Markelius, Connor Wright, Joahna Kuiper, Natalie Delille, and Yu-Ting Kuo. 2024. The mechanisms of AI hype and its planetary and social costs. *AI Ethics* 4, 3 (Aug. 2024), 727–742.

[22] Reza Hadi Mogavi, Chao Deng, Justin Juho Kim, Pengyuan Zhou, Young D Kwon, Ahmed Hosny Saleh Metwally, Ahmed Tlili, Simone Bassanelli, Antonio Bucchiarone, Sujit Gujar, et al. 2024. ChatGPT in education: A blessing or a curse? A qualitative study exploring early adopters' utilization and perceptions. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100027.

[23] Hyanghee Park and Daehwan Ahn. 2024. The promise and peril of ChatGPT in higher education: opportunities, challenges, and design implications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.

[24] Romain Puech, Jakub Macina, Julia Chatain, Mrinmaya Sachan, and Manu Kapur. 2025. Towards the Pedagogical Steering of Large Language Models for Tutoring: A Case Study with Modeling Productive Failure. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 26291–26311. doi:10.18653/v1/2025.findings-acl.1348

[25] Mona Sloane, David Danks, and Emanuel Moss. 2024. Tackling AI hyping. *AI Ethics* 4, 3 (Aug. 2024), 669–677.

[26] Julie M. Smith, Jesse Dukes, Josh Sheldon, Manee Ngozi Nnamani, Natasha Esteves, and Justin Reich. 2025. A Guide to AI in Schools: Perspectives for the Perplexed. https://tsl.mit.edu/ai-guidebook/. Retrieved from https://tsl.mit.edu/ai-guidebook/.