

A Multi-Year Investigation of Undergraduates' Appraisals of Generative AI and Fairness in Higher Education

Sunday Stein, Victoria Delaney, Mich Manrique

Department of Mathematics and Statistics

San Diego State University

San Diego, CA, USA

{sstein9540, vldelaney, mmanrique7330}@sdsu.edu

Abstract

Our work examines the patterns in undergraduate students' appraisals of genAI fairness in school curriculum across a two-year horizon. We identify and present differences in students' fairness appraisals and rationales between this academic year and last, delving into differences in the rationales underlying these appraisals and their implications for higher education policy.

1 Introduction

Technologies progress through phases of skepticism, rejection, acceptance, and adoption from the perspective of public consumers (Davis, 1989). The degree and rate at which adoption occurs in particular contexts can be influenced by a number of cultural, ideological, and social factors that are upheld by communities of practice. For instance, when handheld calculators were introduced into public schools in the 1970s, many communities of mathematics teachers initially rejected them due to fear of arithmetic skill atrophy and its downstream effects on students' mathematical reasoning abilities (Kline, 1973). These attitudes and beliefs, in turn, have contributed to ongoing debates (e.g., "The Math Wars;" Schoenfeld, 2004) about what students ought to learn in school and for what purpose, including how students use technologies for learning. Advances in generative AI (genAI) tools continue to fuel similar debates today.

Taking a sociocultural perspective that tools mediate forms of activity that influence how and what students learn (Engeström, 1987), and adopting Subramonyam et al.'s stance that genAI prompting is in part an extension of one's own goals, cognitive abilities, and knowledge of genAI systems (2024), we have sustained interest in the

complex and potentially shifting question of genAI fairness in higher education and its policy implications. Our prior work on this topic suggests that undergraduates' genAI fairness appraisals are dependent on a number of cognitive, epistemological, and techno-normative rationales (Delaney et al., 2025). However, genAI capabilities and the public's beliefs about how genAI tools might be used in learning are continuing to develop in ways that impact education (e.g., Eames et al., 2026; Kinstler, 2026). Additionally, genAI's capabilities have changed dramatically since we began studying it, and collegiate programs and curricula writ large have not. To that end, we wondered the extent to which students appraise genAI use on school assignments as fair, and qualitatively, how patterns in their appraisals and rationales compare across two student samples collected one year apart.

RQ1: What (a) appraisals and (b) rationales do students give about generative AI fairness in the context of their schooling?

RQ2: How do patterns in students' appraisals and rationales compare between the 2025 and 2026 samples?

2 Background and Framing

2.1 Changes in Students' School-based Engagements with GenAI

GenAI has rapidly shifted from an unfamiliar and disruptive technology to a common feature in the undergraduate experience. In a 2025 study conducted by Lumina Foundation-Gallup, over half of surveyed college students reported using GenAI weekly (Marken, 2026). This adoption suggests that genAI is a part of many students' academic experience, in which studies have shown that students use these tools to complete coursework, seek explanations, draft written responses, and navigate disciplinary expectations

(Baek et al., 2024; Hon, 2026; Martín-Moncunill & Alonso Martínez, 2025). Additionally, students may position genAI as a tool that helps reduce barriers for students with learning difficulties, neurodiversity, or language barriers (Gruenhagen et al., 2024). Given these uses, students may position genAI as a valuable resource in their coursework, especially as AI literacy becomes an increasing expectation in the workplace.

The normalization of genAI use among students and instructors is also tied to workforce expectations, AI literacy is considered part of preparation for future work (“Student Generative Artificial Intelligence Survey 2026,” 2026). Lund et al. (2026) similarly position AI literacy as an “essential skill alongside information literacy,” suggesting that students’ ability to use, evaluate, and integrate tools is becoming part of the expectations for career readiness (p. 4). Institutions have also begun to respond to workforce expectations. An early survey of genAI and higher education policy by Moorhouse et al. (2023) explored worldwide AI adoption and found that some higher education institutions have modified their assessment guidelines to integrate genAI rather than prohibit its use. Given the shift in perspective on the relevance of genAI use in future work, students’ appraisals of what constitutes “fair” genAI use in academic settings may be adapting to these new expectations.

2.2 Policy Inconsistency and Instructor-Level Variation

Institutional and classroom-level responses to genAI remain inconsistent (Lau et al., 2026). Amoozadeh et al. (2024) describe how authority is often distributed across departments, programs, and instructors, making it difficult for institutions to coordinate a unified genAI response. Without a primary source of guidance or enforcement, genAI policies can vary widely across instructors and courses, leaving students confused or unsure about the appropriate use of genAI. They also note that even with institution-wide policies, many students report being unaware of their universities’ rules or guidelines on responsible AI use. The speed of technological adaptation further complicates this problem, where instructors must adapt their policy to real-time changes in AI model availability and resources. These inconsistencies may lead to changes in students’ fairness appraisals, as students may be encouraged to use genAI in one course and

discouraged in another. Policy inconsistency also emerged in our prior study, as students reflected on conflicting experiences of instructors encouraging or discouraging genAI use (Delaney et al., 2025).

2.3 Student Appraisals of Fairness

Students’ genAI fairness appraisals are shaped by overlapping concerns about access, academic integrity, and learning. Moorhouse et al. (2023) found that students raised concerns that their peers may have access to more advanced genAI, such as premium versions of ChatGPT, which can affect the fairness of assessment. Baek et al. (2024) similarly highlight concerns around income inequality in genAI use, with students from lower-income backgrounds reporting concerns about the future cost and accessibility of the technology. Other sources of variation in students’ appraisals stem from how much technology integration is acceptable and what forms it takes in learning processes. Johnston et al. (2024) found that students often distinguished between using genAI for limited academic support, such as grammar assistance, and using it to produce entire assignments. Building on this work, we compare two student samples interviewed in different years to examine similarities and differences in students’ fairness appraisals and rationales across the two samples.

3 Methods

3.1 Research Context and Study Design

Our study was conducted at a large university that serves mostly undergraduates in the southwestern United States in March-May 2025 and 2026. We compare two independent samples of undergraduate students interviewed during these two data collection periods. The 2025 sample included 36 students, with additional information about this data collection reported in Delaney et al. (2025). For the 2026 interviews, we recruited 30 undergraduate students from various majors through flyers and in-person announcements for an interview study about how they selected tools, including genAI (any available chatbot or version), to solve math problems. The samples were broadly similar with respect to students’ class standing and genAI use, although the 2026 sample included fewer students in liberal arts and social science degree programs. The 2025 and 2026 data collections were conducted under approved

institutional review board protocols IRB-25-0023 and IRB-26-101, respectively.

We focus on 10-minute opening section of the interview, which included seven questions about students' understandings and uses of genAI (Table 1). The present study specifically examine responses to Question 7, which asked whether students believed using AI tools in school was fair and why. These questions preceded participants' interaction with genAI during the interview and elicited their perspectives on genAI use generally rather than on a particular platform or model.

Item	Interview Question and Probes
1	What is AI in your view?
2	On the survey you took prior to this study, we noted your preferred AI tool is {blank}. Why do you prefer this/these tools? <i>Probe: Do you have a paid membership to any AI tools?</i>
3	Can you describe how you typically use AI tools, if at all, (a) in your math courses, (b) in your coursework, and (c) in your day-to-day life?
4	What kinds of tasks or problems do you find AI tools most helpful for?
5	How do you decide when to use an AI tool versus solving a problem without one? <i>Probe: Do you find that AI influences the way you approach problem solving?</i>
6	Have you encountered any challenges or limitations while using AI tools? If so, what are they? <i>Probe: If AI tools give you an unclear or unexpected response, how do you typically handle that?</i> <i>How (if at all) do you think about potential AI hallucinations when you use it?</i>
7	Do you think using AI tools in school is fair? Why or why not? <i>Probe: To what extent do your professors encourage/discourage AI use for school?</i>

Table 1: Opening interview questions

Each student's interview was recorded with audio and video and initially automatically transcribed. Then one researcher cleaned each transcript by hand, correcting for acronyms, participant cross-talk, and speaker misattribution. The cleaned transcripts were our primary textual data for coding and analysis. For anonymity, we will refer to the students in our study in the order in which they were analyzed, with P-10 denoting the tenth student.

3.2 Data Analysis

Our unit of analysis is each student's conversational turn to the question, "Do you think using AI tools in school is fair?" We use conversational turns because some students answered the question with an immediate clear fairness judgment and rationale, and for other students who responded ambiguously, the interviewer was trained to follow up with clarifying questions until the student gave a fairness judgment. For example, when P-03 initially responded, "Depends on how you use them," the researcher asked, "Can you elaborate a little more?" and continued to engage the student until they gave a rationale for their appraisal. We then coded the data in two phases. In Phase 1, three researchers worked individually, deductively coding all students' appraisals as "Fair," "Mostly Fair," "Depends," "Mostly Unfair," or "Not Fair," in alignment with five categories that emerged in our previous study (Delaney et al., 2025; Saldaña, 2021). We give responses with a "mostly" qualifier when a student's fairness judgment contains either entirely positive or negative rationales, with one of the opposing valence. We assign "depends" when the appraisal is a mixture of positive and negatively-valenced rationales greater than 1:n. There was one student response that did not have a fairness appraisal, which we coded as "Unclear." Our initial weighted Kappa coefficient was 0.825, signifying relatively high agreement. We then discussed each disagreement until reaching consensus on each student's genAI fairness appraisal.

In Phase 2, we coded students' rationales for their fairness appraisals. Each researcher was assigned a subset of 20 student responses by the first author to code independently. We initially applied the rationale codes developed in the previous year's study (Delaney et al., 2025). When a student's rationale did not align with the existing codebook, we developed new codes inductively to capture the meaning of their response (Saldaña, 2021). For instance, we developed the code, "*genAI during assessments*," for the 2026 cohort because one year ago, students were not mentioning using genAI during assessments. After initial coding, we met to compare our rationale codes, discuss discrepancies, and reach agreement on (a) the rationale codes themselves and (b) what qualified as a new rationale. We developed new rationales collectively to ensure that their names

and definitions accurately represented students' fairness appraisals. Because some students provided multiple examples or justifications, individual responses could be assigned more than one rationale code. For example, some students described both fair and unfair use cases of genAI, which led us to assign multiple rationale codes to a single response.

4 Findings

4.1 Student Appraisals on GenAI Fairness

We summarize and compare students' fairness appraisals from 2025-2026 in Figure 1. We were surprised to learn that, despite the additional year of broader genAI adoption and development between data-collection periods, students in the 2026 sample did not tend to rate genAI as more fair than those in the 2025 sample. The proportion of fairness appraisals in every response category remained mostly consistent. In fact, a greater proportion (48.3%) of 2026's sample elaborated that genAI fairness "depended" (as we will see, it was their rationales for fairness that changed). This pattern suggests continuity across samples, where students tend to evaluate the fairness of genAI contingent on how, where, and what purpose it was used even after one year of technological development. Responses coded as "Depends" typically included the phrase "it depends" or described balanced appraisals in which genAI would be fair in some contexts and unfair in others. For example, P-13 distinguished between genAI use on homework assignments and exams:

"I think AI tools are fair for homework assignments or assignments in which you're not being purely tested on what you know on the spot, like in your mind, but I think it'd be unfair to incorporate AI into exams, because then you wouldn't really get to gauge the person who you are assessing."

P-13's response illustrates how students' fairness appraisals may be tied to specific practices that they deem to be fair or unfair, which we will explore in the next section. They also implicitly suggest a difference between assessments in-class ("on the spot") which may be designed to measure only the knowledge of the student, versus assignments outside of class, which may be designed for other purposes (e.g., to give students practice toward a skill).

Undergraduate Appraisals of AI Fairness by Interview Year




Figure 1: Undergraduate students' generative AI fairness appraisals by interview year.

Students' appraisals coded as Mostly Fair/Unfair reflected a general sentiment toward one judgment, with exceptions. For example, P-29's reflection of Mostly Fair excluded tests but appraised genAI as fair for learning on anything else:

"I think there's a point where [if] you're using it to cheat on tests, that's not good...no one should be doing that. But I think as a tool to help in further learning...I think that's where the world is going...and using it to help with problems...you're struggling on."

P-29's response reflects an ethical stance that tools should be used for learning purposes, especially when they are normalized ("that's where the world is going,"), and that using genAI on tests represents a breach of tool use for learning. Students from 2025 had very similar perspectives, although they tended to be more specific about the exceptions to intellectual work that genAI should or should not do on homework assignments. For instance, a 2025 participant remarked that "mostly fair" use on a writing assignment might involve genAI as an editor, not as the primary author. While the distribution of appraisal categories remained broadly similar across samples, the 2026 responses continued to show that students' appraisals are not absolute, and even participants who leaned fair or unfair frequently identified exceptions.

4.2 Comparison of Student Appraisals on GenAI Fairness

Figure 2 displays the frequency of rationale-code occurrences by appraisal category. Because students could offer multiple rationales in a single response, the counts in Figure 2 represent coded rationales rather than unique students. Lettered rationale codes were developed in our previous study, while asterisked codes emerged during the

2026 interviews and analysis. Across rationale categories, responses coded as “Depends” appeared frequently, indicating that students often justified their fairness appraisals by identifying contexts in which genAI use could be fair and other contexts in which it could be unfair.

Although students’ appraisal categories were broadly similar across samples, the rationales observed in the two samples differed in several ways. In the 2025 data, several of the most frequent rationales positioned genAI as an aid for clarification, scaffolding, or co-authoring with evidence of students’ or genAI’s contributions. In contrast, the 2026 data included more rationales that explicitly framed genAI as a tool for learning in general, rather than providing examples of how genAI can be considered fair during learning. The 2026 sample also included more rationales that focused on the conditions under which genAI use undermines learning, assessment, or preparation for future professional practice. This suggests that students’ appraisals remained largely conditional, while the substance of their rationales became more varied and explicitly tied to learning and assessment.

One notable shift from 2025-2026 was the emergence of *cognitive offloading* for assignments as a frequent rationale. This code captured responses in which students expressed concern that genAI allowed students to completely bypass thinking, effort, or problem-solving associated with schoolwork. For example, P-03 framed fairness as context-dependent, claiming, “If you’re using [genAI] to write an entire essay or do all your math problems, then it’s not really fair to people who are actually putting in the work.” This response reflects a concern that genAI becomes unfair when it substitutes for student cognitive work rather than supporting their learning. In contrast, students in the 2025 sample were more likely to report that using genAI to tutor them on math problems, give additional examples of math problems, or to help them to brainstorm ideas for an essay were acceptable genAI uses. None in 2025 mentioned cognitive offloading, perhaps because they explicitly stated forms of cognition and work that they might directly contribute.

Students in 2026 also raised concerns about genAI fairness in relation to future professional competence. The *workplace deprofessionalism* code captures responses in which students worried that too much genAI in coursework could leave

students underprepared for fields involving human safety or embodied expertise. For example, P-13 explained, “if you’re studying something [...] for example, nursing, I feel like you can’t be dependent on [genAI] because [...] you’re studying to help a human being.”

Similarly, P-08 described a hypothetical student training to become a surgeon, who may rely on genAI to decide what to do “mid-surgery.” These responses suggest that some students evaluated fairness in relation to whether academic work with genAI prepares students for responsible professional practices that directly involve human-to-human encounters.




Figure 2: Rationale-code occurrences for students’ genAI fairness appraisals by year and appraisal category

In contrast to the previous concern-oriented rationales, students in 2026 continued to describe genAI as a form of academic support. Several rationales from the 2026 sample positioned genAI as useful for searching for information (P-09), supporting learning (many students), functioning like a tutor (P-23), or addressing limitations in human instructional support (P-24, P-29). In these responses, students described genAI as a resource that could provide explanations, immediate

feedback, or assistance when instructors were unavailable due to time and class-size restraints. The theme of genAI as a form of academic support also emerged in the 2025 data, where students frequently discussed using genAI to clarify or scaffold assignments, or to use it as a co-author with proper citations. P-04, a Japanese major, identified potential benefits of genAI to neurodivergent students. Similar to findings reported by Gruenhagen et al. (2024), P-04 suggested that students on the autism spectrum may benefit from genAI as one way to ask as many questions as needed, at any time of day, without relying on a human instructor's availability or attention. These responses suggest that students may view genAI as a means of making academic participation more accessible in universities.

5 Discussion

Our findings differed from our initial conjecture. We hypothesized that, based in part on Davis's model of technology acceptance (1989), students would be more inclined to appraise genAI for school assignments as "Fair (Yes)" or "Mostly Fair (Yes/It depends)" because students and instructors have now had multiple years to grow accustomed to its features and adapt accordingly. Instead, we observed that students still held ambiguity about genAI fairness, and the focus of fairness appraisals differed across samples, from what to do with genAI within an assignment (2025) to larger grain-sized features of education, such as tests, the workforce, and societal expectations.

Why might these differences across samples have occurred? We offer several possible conjectures. One, 2026 was largely concentrated on innovations in agentic genAI technologies; products such as Codex that are designed to accomplish end-to-end workflows. No student brought up that they preferred to use agentic genAI tools, although P-10 (a computer science major) had noted: "I'm trying to use AI less, I guess I'm trying to learn more of the [programming] concepts and apply myself more, just because I'm a third year and trying to actually apply myself." It is possible that some undergraduates realize that today's most novel agentic genAI tools disrupt traditional goals of undergraduate education to the extent that they may undermine the purpose of their degree (Labaree, 1997). Hence, in the data, we may continue to see "it depends" responses from students because genAI developments may lead

them to continually recalibrate how they use such tools in ways that still lead to learning.

Another possible reason is that students' uses of genAI tools may have begun to cross a threshold of traditional conceptions of fairness in 2026. In 2025, students' rationales tended to illustrate ambiguity on how to use genAI during assignments; however, there were plenty of technologies before genAI that students could use to "support" them on assignments that were ethically ambiguous from the perspective of instructors. Chegg.com, for instance, branded itself as an "innovative learning platform" but was renowned for distributing answer keys to textbook exercises. In 2026, students' rationales instead reflected a consistent "fair as long as not cheating on tests" sentiment. This begs the question of whether students are beginning to observe genAI cheating on tests, a form of activity that has historically led to greater negative consequences in education than using ambiguous forms of support on homework assignments. Assessments in higher education generally comprise a larger portion of students' grades; another reason why students may be flagging this particular issue as unfair.

5.1 Policy Implications, Limitations, and Concluding Provocations

Prior work, including ours, has suggested that genAI use policies at the course level can espouse instructors' goals and beliefs while potentially confusing students and posing additional hardships about whose genAI policy they should adhere to at any given time (Johnston et al., 2024; Lau et al., 2026). Many institutions have begun to adopt genAI policies at the department, college, and institution levels in response. Yet, as our findings show, students may still rely on their own appraisals about genAI fairness, and very few students in our samples across both years referenced course or department policies at all. Analyzing this data has encouraged our team to step back and ask ourselves, "What is a genAI policy intended to do in an increasingly genAI saturated world?" Given that 2026's sample showed increased polarized appraisals of AI fairness in the classroom compared with 2025, how might educators and administrators respond?

We admit that we do not yet know the answers to these questions. At the time of this writing the United States is experiencing a visible, public backlash to AI and data centers (Goldmacher,

2026), so it is unclear to us if these tensions follow hypothesized patterns in Davis' technology adoption curves or something more contentious. As a starting point, we contend that students' perspectives can at least help stakeholders within higher education have critical conversations about what kind of genAI measurements and/or student data they are collecting, and what sorts of policy decisions are made with such data. One affordance of our work is that it garnered a variety of new genAI rationales, some of which we were not expecting, as interview studies and focus groups often do (Johnston et al, 2024). We encourage administrators to weigh the measurement affordances of various tools as they ask themselves what they want to learn about students and genAI, and how the data collected will be used to make decisions or craft policy. For instance, administrators might consider whether to send the same student AI use survey every year (to measure consistency or compare patterns across years) or to update survey questions in response to changes motivated by genAI tools (to expand what is captured by their measures). They might also consider whether surveys are the best way to capture large-scale data about genAI interactions and fairness, or if other modalities (e.g., AI-administered interviews) could yield different or complementary insights to surveys.

Despite most students' inattention to genAI course policies, we still believe that higher education instructors must be given time, space, and intellectual freedom over the curriculum-tool relationship in their courses (Brown, 2009; Delaney & Lee, 2026). However, our data suggest that undergraduates might benefit from multiple reminders throughout the semester about course genAI policies and why they matter/are directly related to course learning outcomes. Degree programs and policies might encourage students to develop human-specific forms of expertise (Scardamalia & Bereiter, 2020; Wu et al., 2025). For example, mathematics teacher educators might spend fewer hours preparing pre-service teachers in mathematics content and more hours developing care-based work with human students.

As we conclude this work, there are limitations that impact the general applicability of our findings. Our small sample traded the statistical power and generalizability afforded by a large sample for rich descriptions of what genAI fairness looks like for students on a day-to-day

basis. Because the 2025 and 2026 data were collected from independent samples, differences observed between years may also reflect variation in the composition or experiences of the two samples intermixed with temporal changes in undergraduate students' fairness appraisals. Furthermore, our study was situated at a large undergraduate institution whose administration openly supported genAI integration, which may have impacted students' fairness appraisals. Other institutions, such as small liberal arts colleges, may have produced a sample with different (more genAI unfair) results because of the thinking and reasoning traditions their curricula and pedagogy embrace. In our future work, we will attempt to expand our sample across institution types and data collection modalities to encompass a broader range of genAI fairness appraisals and the factors that underpin them.

References

- Amoozadeh, M., Daniels, D., Nam, D., Kumar, A., Chen, S., Hilton, M., Srinivasa Ragavan, S., & Alipour, M. A. (2024). Trust in Generative AI among Students: An exploratory study. *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, 67–73. <https://doi.org/10.1145/3626252.3630842>.
- Baek, C., Tate, T., & Warschauer, M. (2024). "ChatGPT seems too good to be true": College students' use and perceptions of generative AI. *Computers and Education: Artificial Intelligence*, 7, 100294. <https://doi.org/10.1016/j.caeai.2024.100294>.
- Brown, M. W. (2009). The teacher-tool relationship: Theorizing the design and use of curriculum materials. In J. T. Remillard, B. A. Herbel-Eisenmann, & G. M. Lloyd (Eds.), *Mathematics teachers at work: Connecting curriculum materials and classroom instruction* (pp. 17–36). Routledge. <https://doi.org/10.4324/9780203884645-11>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>.
- Delaney, V., & Lee, V. R. (2026). Building students' active participation from teacher knowledge and adaptation. In Z. A. Traga Philippakos, E. Katsarou, & K. Sipitanos (Eds.), *Artificial intelligence in literacy education: Foundations, practices, and innovations* (pp. 70–84). Routledge. <https://doi.org/10.4324/9781003661443-6>

- Delaney, V., Stein, S., Sawi, L., & Hernandez Holliday, K. (2025). Undergraduate students' appraisals and rationales of AI fairness in higher education. In J. Wilson, C. Ormerod, & M. Beiting Parrish (Eds.), *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full papers* (pp. 329–336). National Council on Measurement in Education. <https://aclanthology.org/2025.aimecon-main.35/>
- Eames, T., Brunskill, E., Yamkovenko, B., Weatherholtz, K., & Oreopoulos, P. (2026). Computer-assisted learning in the real world: How Khan Academy influences student math learning. *PNAS*, 123(1), e2507708123. <https://doi.org/10.1073/pnas.2507708123>.
- Engeström, Y. (1987). *Learning by expanding: An activity-theoretical approach to developmental research*. Orienta-Konsultit.
- Goldmacher, S. (2026, August 23). The data center backlash bursts into the midterms. *The New York Times*. <https://www.nytimes.com/2026/08/23/us/politics/data-centers-midterm-elections.html>
- Gruenhagen, J. H., Sinclair, P. M., Carroll, J.-A., Baker, P. R. A., Wilson, A., & Demant, D. (2024). The rapid rise of generative AI and its implications for academic integrity: Students' perceptions and use of chatbots for assistance with assessments. *Computers and Education: Artificial Intelligence*, 7, 100273. <https://doi.org/10.1016/j.caeai.2024.100273>.
- Hon, K. (Kai L. (2026). Generative AI in Higher Education: A Systematic Review of Its Effects on Learning Outcomes and Academic Performance. *Journal of Educational Technology Systems*, 54(3), 537–560. <https://doi.org/10.1177/00472395251400089>.
- Johnston, H., Wells, R. F., Shanks, E. M., Boey, T., & Parsons, B. N. (2024). Student perspectives on the use of generative artificial intelligence technologies in higher education. *International Journal for Educational Integrity*, 20(1), 2. <https://doi.org/10.1007/s40979-024-00149-4>.
- Kinstler, L. (2026, June 1). A university system went all in on A.I. Now it's tearing itself apart. *The New York Times Magazine*. <https://www.nytimes.com/2026/06/01/magazine/ai-university-college-california.html>.
- Kline, M. (1973). *Why Johnny can't add: The failure of the new math*. St. Martin's Press.
- Labaree, D. F. (1997). Public Goods, Private Goods: The American Struggle Over Educational Goals. *American Educational Research Journal*, 34(1), 39–81.
- Lau, S., Boroojeni, K., Keeling, H., & Marroquin, J. (2026). Barriers that programming instructors face while performing emergency pedagogical design to shape student-AI interactions with generative AI tools. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (CHI '26). ACM. <https://doi.org/10.1145/3772318.3790682>.
- Lund, B. D., Warren, S. J., & Teel, Z. A. (2026). *Measuring University Students Satisfaction with Traditional Search Engines and Generative AI Tools as Information Sources* (Version 1). arXiv. <https://doi.org/10.48550/ARXIV.2601.00493>.
- Marken, S. (2026, April 2). *AI Is Routine for College Students, Despite Campus Limits*. Gallup.Com. <https://news.gallup.com/poll/704090/routine-college-students-despite-campus-limits.aspx>.
- Martín-Moncunill, D., & Alonso Martínez, D. (2025). Students' Trust in AI and Their Verification Strategies: A Case Study at Camilo José Cela University. *Education Sciences*, 15(10), 1307. <https://doi.org/10.3390/educsci15101307>.
- Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open*, 5, 100151. <https://doi.org/10.1016/j.caeo.2023.100151>.
- Saldaña, J. (2021). *The coding manual for qualitative researchers* (4E [Fourth edition]). SAGE Publishing Inc.
- Schoenfeld, A. H. (2004). The math wars. *Educational Policy*, 18(1), 253–286. <https://doi.org/10.1177/0895904803260042>.
- Student Generative Artificial Intelligence Survey 2026. (2026, March 12). *HEPI*. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2026/>.
- Scardamalia, M. & Bereiter, C. (2020), Will knowledge building remain uniquely human? *QWERTY-Interdisciplinary Journal of Technology, Culture and Education*, 15(2), 12-26.
- Subramonyam, H., Pea, R., Pondoc, C., Agrawala, M., & Seifert, C. (2024, May). Bridging the gulf of envisioning: Cognitive challenges in prompt based interactions with LLMs. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1–19). <https://doi.org/10.1145/3613904.3642754>
- Wu, J. Y., Lee, Y. H., Chai, C. S., & Tsai, C. C. (2025). Strengthening human epistemic agency in the symbiotic learning partnership with generative artificial intelligence. *Educational Researcher*, 54(6), 358–368. <http://doi.org/10.3102/0013189X251333628>