

Participatory, not Punitive: Student-Driven AI Policy Recommendations in a Design Classroom

Kaoru Seki*

kaoru2020@umbc.edu

Univ of Maryland, Baltimore County
Baltimore, MD, USA

Manisha Vijay*

mvijay1@umbc.edu

Univ of Maryland, Baltimore County
Baltimore, MD, USA

Yasmine Kotturi

kotturi@umbc.edu

Univ of Maryland, Baltimore County
Baltimore, MD, USA



Figure 1: Ten graduate design students at a minority-serving university co-authored 10 policy recommendations through student-led workshops and follow-up interviews, then visualized them in a zine—a shareable, DIY booklet. Student-led discussions revealed fresh perspectives on existing policies and surfaced considerations often missed in top-down AI mandates. The initial zine pages shown above, and the full zine is viewable at <https://ykotturi.github.io/zines/>

Abstract

Generative AI is reshaping education, yet most university AI policies are written without students and focus on penalizing misuse. This top-down approach sidelines those most affected from decisions that shape their everyday learning, resulting in confusion and fear about acceptable use. We examine how participatory, student-driven AI policy design can address this disconnect. We report on a three-part workshop series in a graduate design course at a minority-serving university in the U.S., where two student leaders facilitated discussions without faculty present. Eight participants shared candid accounts of their AI use, co-authored ten policy recommendations, and visualized them in a zine that circulated across campus. The resulting policies surfaced concerns absent from top-down governance, such as the double standard of requiring students to disclose or abstain from AI use while faculty face no such expectations. We argue that engaging students in AI governance carries value beyond the resulting policies, and offer transferable strategies for fostering participation across disciplines—a model for calling students *in* rather than calling students *out*.

*Both authors contributed equally to this research.

Keywords

generative AI, participatory governance, higher education, student perspectives, design pedagogy

ACM Reference Format:

Kaoru Seki, Manisha Vijay, and Yasmine Kotturi. 2026. Participatory, not Punitive: Student-Driven AI Policy Recommendations in a Design Classroom. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain*. ACM, New York, NY, USA, 29 pages. <https://doi.org/10.1145/3772318.3790691>

1 Introduction

Generative artificial intelligence (AI) is rapidly reshaping teaching and learning practices across universities. Although not the first technology to disrupt higher education [79], generative AI technologies capable of producing text, images, and other media [111] have gripped administrators, faculty, and students alike, all racing to manage its challenges and opportunities [115]. Challenges are far-reaching: from questions of students' academic integrity [18] and over-reliance [137], to reductions in independent learning and weakened skill transfer [72]. Overwhelmingly, responses to these challenges have led to university policies that emphasize the prevention of misconduct and plagiarism [12], which are often framed in punitive terms rather than pedagogical support [43, 76].

Across universities nationwide, one pattern holds in AI policy discourse: students' perspectives are largely absent [24, p.134]. And yet, students are the primary group affected by universities' AI policies [101, 136]. As noted in recent calls for a "participatory turn in

AI" [37], omitting key stakeholders' perspectives and experiences from governing decisions lessens the efficacy of these decisions, and may even result in harmful outcomes due to exclusionary thinking [44]. Students bring nuanced perspectives into how AI is *actually* used in everyday coursework—insights often invisible in top-down policy discourse and enforcement [68]. For instance, students are enthusiastic about using AI to improve personalized learning opportunities, self-reflection, and overall efficiency [63]. In addition, students are concerned with AI extending instructors' surveillance capabilities [49], fueling incorrect accusations of AI plagiarism [47], and eroding trust in student-teacher relationships [51]. Despite the richness and importance of these perspectives, there is little evidence that they are taken into account when shaping generative AI policy [133].

Therefore, in this paper, we take up calls to incorporate student perspectives in AI policies [24, 29, 51, 63, 101]. We frame students as "lead users" [130], or early adopters of AI technologies. In other words, we position students as experts in their own lived experiences with generative AI—offering insights that faculty and administrators can ultimately learn from when determining how (and how not) to integrate these technologies into the classroom. In doing so, we work towards student-driven policy recommendations, where all stakeholders' perspectives, especially those most affected by the policies, should have the ability to mold such policies [68].

We conducted a three-part, participatory design workshop series with follow-up interviews with 10 students from a graduate-level design studio course at University of Maryland, Baltimore County (UMBC): a minority-serving public institution in the mid-Atlantic region of the United States. To address inherent power dynamics between students and faculty, we implemented specific procedural safeguards: workshop activities were student-led to foster open dialogue without faculty oversight, and the research team ensured that faculty had no access to raw transcripts. By grounding workshops in a class that all participants had recently taken, discussions and recommendations were based on recent experience [68], making for higher quality recommendations as students do not typically have prior policy writing experience [81].

After an iterative design cycle (brainstorming, drafting, applying, and refining policies), we report on 10 resulting policies (See Figure 2). Policies were embedded in a student-authored zine—a DIY booklet [38, 74] used as an artifact and method to promulgate information—which was circulated widely to support student-driven discourse across the campus (See Figure 1), as well as archived in UMBC's Library¹. We view this distribution strategy as a form of participatory infrastructuring [34]: by archiving the zine and circulating it through physical and digital channels, we moved beyond the workshop series as a singular event, instead creating lasting resources that allow student voices to persist in institutional memory and continue shaping policy long after the initial workshop activities concluded.

In addition, through thematic analysis and triangulation of workshop dialogue, interviews, and artifacts (e.g., zines, design activities), we find how participants, when provided the environment to speak candidly, gradually disclosed more layered accounts of AI

use, such as when and why they used AI that contradicted existing course policies. While student-authored policy recommendations reflect students' leniency towards AI use writ large [103, 106, 112], they also capture the nuanced, lived perspectives of those most directly implicated in governance decisions. Importantly, we observed how the process of authoring policies shifted participants' practice towards more intentional, reflective AI use, suggesting that the process of participation may be equally important to the outputs.

Together, this paper makes three contributions. First, we contribute an empirical account of students' candid AI practices in a graduate design studio—including misuse, policy gray areas, and uneven enforcement. Second, we contribute a transferable model for student-driven AI policy design that combines faculty-free workshops, student-led facilitation, and zine-based participatory infrastructuring to lower barriers to candid expression and support participatory governance in higher education. Third, we contribute ten actionable, design-oriented policy recommendations authored by students, situated within the broader landscape of institutional AI policies—highlighting where student perspectives align with, extend, or diverge from prevailing administrative approaches. We reflect on both the challenges and the benefits of student-driven policy recommendations, arguing that the participatory process itself—calling students in as co-authors of governance rather than subjects of regulation—carries value beyond the policies it produces.

2 Related Work

We draw on three bodies of scholarship to inform this work: how generative AI is (and is not [96]) reshaping the classroom in higher education, how diverse participation improves efficacy and trust in AI governance, and how participatory infrastructuring enacted through zine-making offers a framework to actualize participatory aims of universities' AI governance.

2.1 Generative AI in the (Design) Classroom

Generative AI has quickly become a contested presence in higher education, continuing a long lineage of technological intervention in the classroom—from calculators in the 1970s to MOOCs in the 2010s—that have disrupted teaching and learning. Unlike earlier technologies, however, the scale and scope of automation possible with generative AI uproots and nullifies many of the safeguards educators rely on to ensure adequate learning [76]. For instance, use of text-generation AI tools (such as ChatGPT or Claude), can reduce students' opportunities to experience "desirable difficulties"—or the productive cognitive struggle that causes long-term retention and skill transfer [19]. As a result of this decreased cognitive engagement, students who use generative AI technologies in unstructured ways may not develop essential skills: critical thinking, problem-solving, and creativity [45, 64]. Given these risks, universities have been compelled to regulate their students' use of AI technologies [61, 76], resulting in a cacophony of regulatory attempts [89].

To understand this capricious policy environment, recent scholarship categorizes institutional responses along a spectrum from restrictive to experimental [119]. On the restrictive end, some universities have implemented a blanket ban on AI use, like Sciences Po

¹A digital version of the zine archived by the UMBC Library can be found here: <https://lib.guides.umbc.edu/c.php?g=1475372&p=11187429>

1. Instructions Instructors should include guidelines with concrete examples of acceptable and not acceptable use of AI for each assignment.	6. Citing AI Use (or Not) It is important for us to credit AI where credit is due. However, sharing chat logs, as many instructors currently require, is tedious and ineffective for both students and graders. Instead, we should share a 2–3 sentence summary for each submission describing how and why we used AI.
2. Ownership We should only use AI for up to 50% of our work on any given assignment, so that the majority reflects our own ideas and effort	7. Hypocrisy In Faculty Use If instructors expect us to be transparent about our use of generative AI, we expect the same transparency from them when it comes to how they use AI in teaching, grading, or creating assignments.
3. Divergent Thinking When using AI for brainstorming, we should push ourselves to explore alternatives, surprising directions, or ideas that feel more personal and meaningful to us	8. English Learners For those of us who are English learners, we should be encouraged to use AI to support our English proficiency in writing by asking for refinement of our text. Also, we should ask for an explanation of the refinements made to further our language proficiency.
4. Job Skills We should be provided opportunities to use AI in ways that reflect how AI is used in real workplaces through coursework. Instructors should stay updated on how AI is used (and regulated) in industry.	9. Feedback We can use AI to help us better understand and respond to peers' feedback, especially when revising our work.
5. Bias We should check if AI-generated ideas include any stereotypes or biased assumptions, such as by asking if any perspective or voice is missing in the response received.	10. Equity Everyone in class should have access to and use the same AI tools or models for each assignment to ensure fairness.

Figure 2: Ten student-drive AI policy recommendations derived from our three-part design workshop series

in Paris, France, which prohibits generative AI use without explicit authorization under threat of expulsion [59]. On the experimental end, initiatives like the University of Florida's "AI Across the Curriculum" [116] and Ohio State's "AI Fluency" initiative [93], claim that all of their graduates (not just their computing and engineering students) will be trained with "essential AI skills" such as using "AI tools to accomplish specific goals in the field of study, and critically assess outputs for accuracy and relevance to the task" [93, 116]. Arizona State University's recent partnership with OpenAI takes this experimentation a step further to launch "innovation challenges" inviting students to co-create use cases for what AI use may be beneficial [105]. And of course, there are many institutions in between these two extremes, such as at University of Maryland, Baltimore County (UMBC)—the site of this study—where recent AI policy guidelines recommend that instructors update their syllabi to signal to students what level of AI use is permitted in class: "Green light" (unrestricted), "Yellow light" (some restrictions), and "Red light" (prohibited).

Across these myriad approaches to regulate student use of AI, from the lenient to punitive, key problems emerge. First, there is often little explanation for the rationale behind the policies. Implementing AI policy without articulating the pedagogical reasoning, whether concerns around learning outcomes, academic integrity, and skill developments, may lead students to see it as arbitrary practices of authority rather than thoughtful pedagogical decisions [117]. Without meaningful rationale, students are less likely to internalize the policy's value, resulting in diminished engagement, autonomous motivation, and performance [9, 36].

Another key challenge with current approaches to AI governance in higher education is a lack of reciprocity: while students must disclose their usage and/or are prohibited from using AI, instructors are not held to the same standard, and can use AI for teaching or grading without similar disclosure rules [77]. This asymmetry reinforces power imbalances and reduces trust between students and faculty [77]. Finally, another challenge includes the enforcement of policies: making sure that students only use AI within

the bounds of allowed use. The predominant approach to address this challenge comes laden with additional problems: AI detection tools, though intended to prevent misconduct, introduce serious risks of false accusations [47, 49]. Unlike plagiarism systems such as TurnItIn.com that compare work against known sources [124], AI detectors infer authorship probabilistically, making accusations far more uncertain and contestable [40]. These issues suggest a fundamental limitation of top-down AI policy design that excludes students—the most affected stakeholders—from its implementation.

2.1.1 Generative AI in the Design Studio: Opportunities and Risks. While generative AI is reshaping higher education broadly, this paper focuses on its impact on design pedagogy. Design studios are a salient site for study because they rely on open-ended inquiry, deliberate practice, and iterative processes [73]—all practices permeable to AI’s multimodal capabilities. Empirical studies paint a complex picture of AI in design which includes both opportunities and risks towards opportunity: generative AI can support early phases of conceptual design—problem framing and ideation [31]. Design students describe large language models (LLMs) as a “second mind” for externalizing and iterating ideas [132], aligning with views of AI as collaborator rather than tool [109]. Additionally, hybrid systems merging physical prototyping with AI open new avenues for low-fidelity exploration [135], and antagonistic AI can provoke divergence rather than premature convergence during ideation [28, 75].

On the other hand, risks persist: AI-assisted ideation can reduce variation of ideas [131], as prompting practices can homogenize outputs when users converge on popular modifiers [94]. This homogeneity is then exacerbated by generative models’ statistical regularities in training data, which steer creators toward shared stylistic outcomes [94, 131]. In other words, even with careful prompting techniques that anticipate AI’s homogeneous tendencies, the underlying technology that powers generative AI technologies can undermine these end-user attempts to course correct. In addition, issues of consent, credit, and compensation remain unresolved: AI models are trained on creatives’ works without consent [69], spurring calls for alternative ownership paradigms [100]. Our study contributes by examining how design students themselves articulate these tensions and propose ways to mitigate risks and pursue opportunities. Given that students are rarely asked to be involved in such governance discussions, such an investigation requires a participatory approach, as discussed in the next section.

2.2 The Participatory Turn in AI

Recent HCI scholarship has focused on the “participatory turn” in AI, where decision-making for AI systems’ design and governance should not be limited to the creators, but also includes those who are impacted by such technologies [37]. This repositioning emphasizes the importance of acknowledging stakeholders’ lived experience and involving them not only in giving feedback on systems but also in co-setting agendas and informing what is built in the first place [17, 44]. In addition to improved efficacy, participatory approaches to AI governance may also serve a trust-repair function. For instance, in higher education where AI infiltration has damaged student-faculty relationships [49, 98], participatory approaches to AI governance can help to repair breakdowns by calling students in

rather than calling students *out*. Students-as-partners scholarship shows that co-creation shifts classroom dynamics from hierarchical to collaborative, with benefits for both student engagement and learning outcomes [23, 86].

Importantly, participatory AI can slide into tokenism without intentional steps to anticipate and course correct [17]. Delgado et al. highlight how many participatory AI efforts remain consultative rather than empowering [37]. Pushing past tokenism requires moving participation upstream, sharing decision authority, and creating safe conditions for dissent. To do this successfully requires a clear-eyed understanding of the barriers to participatory design approaches: it is time and resource-intensive, may not scale easily across institutional contexts, and risks placing burdens on participants asked to perform tasks without formal training [20, 129]. In the context of participatory governance in the classroom, tools like PolicyCraft [68] mitigate these challenges with “case-grounded deliberation,” structuring policy discussion around concrete scenarios, seeded by facilitators, to make governance decisions easier for non-experts, such as students.

In this paper, we consider another approach to fostering student participation in AI governance, while acknowledging the challenges that come alongside fostering participation. We consider power asymmetries between students and faculty, even in ostensibly participatory spaces, requiring deliberate structural safeguards [22]. We frame students as “lead users”—or early-adopters [130]—who have particular expertise from lived experience that instructors may lack, and design power-aware, student-led workshop activities in order to lower barriers to candid expression and resist tokenistic consultation. Finally, we center student discourse around a zine-making activity in order to emphasize not only the resulting artifact, but also the importance of deliberation among students along the way (as described in the next section).

2.3 Participatory Infrastructuring through Zine-Making

In this work, we frame zines—small, self-published booklets [15, 38]—as infrastructure: material media that circulate student perspectives across classrooms and institutions while catalyzing social formations and shared knowledge. Their low-cost, DIY format positions them as alternative media for grassroots circulation [85], amenable to candid authorship and resistant to polish and performativity. Zine-making has played a critical role in feminist, punk, and activist communities to amplify marginalized voices and foster counterpublics [38, 74]. More recently, zines have been used as a participatory design method in HCI to surface marginalized perspectives [52, 54]. In contrast to top-down policy memos that reduce student agency to checkbox compliance [46], zines may allow students to narrate their own experiences, reflect on trade-offs, and suggest alternatives. It is important to note that zine-making can be challenging for those unfamiliar with their format, but strategies such as detailed explanations, examples, and templates can ameliorate this [54].

We view zine making through the lens of participatory infrastructuring, which emphasizes creating sociotechnical conditions that allow publics to form—not pre-defined stakeholder groups, but emergent collectives constituted through shared concern and active

engagement [34]. These publics participate not merely as respondents, but as co-constructors of the problem space itself [14, 34]. Through a participatory infrastructuring lens, we frame student-led zine-making as collective world-building: students shifting from a fixed stakeholder group to a reflexive public capable of intervening in institutional discourse. Student-centered approaches to policy articulation are especially important in higher education, where classroom authority is traditionally asymmetrical. Scholars of critical pedagogy argue that shifting power toward students can improve learning outcomes and foster civic capacities [56, 90, 113]. Student-driven policies may be more likely to be respected and internalized when they reflect lived realities [24]; recent work by Pu et al. demonstrates the value of engaging students in co-designing AI guidelines grounded in their concerns [101].

3 Methods

We draw on community-based participatory design [52] and conducted a three-part, in-person design workshop series to facilitate policy articulation, application, and iteration. This workshop series was led by two student research assistants, who are joint first authors of this paper.

3.1 University setting

This study was conducted at University of Maryland, Baltimore County (UMBC): a minority-serving, public university in the U.S. mid-Atlantic. To contextualize AI use among UMBC's student body relative to other universities, a campus survey indicated slightly slower AI uptake relative to national reports: 27% reported using AI in coursework (vs. 42% in a national study [125]); 73% had not taken a formal AI course and AI use skewed heavily towards STEM majors [42, 82]. Therefore, UMBC represents a context where formal AI policies and institutional support structures are still early-stage.

3.2 Participants and recruitment

Participants were recruited from the last author's graduate design studio course, HCC 629: Fundamentals of Human-Centered Computing, taught the preceding semester before the study (Fall 2024). This is a required course for the master's and Ph.D. students in the department's HCI specialization, and focuses on fundamental design principles and human factor concepts to guide effective user interface design. Of 21 enrolled students, 10 completed a screener with questions such as "Describe how you used genAI in HCC 629?" and "At UMBC, how do existing AI policies support/hinder your learning?" Eight students were selected to vary course performance [62], AI use, schedules, and degrees (e.g., HCI, information science, software engineering). Informed consent from all participants was acquired, and all ethical procedures adhered to the standards of the university institutional review board, as per the approved IRB protocol by which this research was conducted. Workshop participants were compensated \$20/hour (11 hours; \$220 total for each workshop participant for the three-part workshop participation). Following the workshop series, we conducted five follow-up interviews: three with workshop participants and two with students who declined workshop participation due to dissenting views on AI or scheduling conflicts. Workshop participants who completed follow-up interviews were compensated an additional

\$20; interview-only participants were compensated \$20 for one hour of their time (described in detail in Section 3.4). In total, 10 students contributed: 8 workshop participants and 2 interview-only participants (degrees: 7 HCI, 2 Information Science, 1 Software Engineering; 9 master's, 1 Ph.D.). To protect identities given public zine authorship, we report only these high-level demographics.

3.3 Three-part workshop series

A total of three workshops were held on Fridays: March 28, April 11, and April 25, 2025, from 12:00–3:00pm; lunch was provided. See Figure 3 for an overview of workshop structure and duration. Each workshop was held in a design studio classroom with modular tables, chairs, and whiteboards to support collaborative activities. Prior to the workshop series, participants completed a pre-workshop survey to capture their motivations to participate, and initial opinions, skills, and uses of generative AI in the graduate design studio course (and beyond). Students' top two motivations to participate included learning how to use AI more ethically, and having their perspectives taken into consideration in AI policy creation.

3.3.1 Workshop 1: Policy Drafting through Candid Conversations. Workshop 1 focused on supporting candid discussions among students as a launching point to brainstorm policy recommendations. Given the sometimes punitive environment surrounding generative AI use, we took steps to create an environment where students felt comfortable disclosing use, voicing concerns, and sharing guidance with each other. First, to ensure participants had a safe space to converse frankly and honestly, these conversations were driven by the two student leaders sans faculty. These two student leaders were graduate peers of the participants rather than instructors or teaching assistants; they held no formal evaluative role (e.g., grading, advising, or assessment responsibilities) in relation to the students who took part in the workshops or follow-up interviews. Second, the faculty PI did not have access to the recordings nor the raw transcript, and could only view a de-identified version of the transcripts (this was explained repeatedly to participants throughout the workshop). Finally, student leaders reminded participants they have essential lived experience with AI tools, and that the study aimed to learn from them (See Supplemental Materials for workshop slides and all other workshop materials).

Think-Pair-Share to Ground Discussions of AI Use. During the first workshop, participants engaged in a "Think-Pair-Share" [78] where they reflected on their use of AI (e.g., how their use fell into a "gray area" of acceptability or how they used AI in ways that violated their courses' AI policy), compared these experiences with a peer, and then partook in a group-wide discussion that elevated key themes. Students were encouraged to ground their reflections in recent experience. To help students situate their experiences within a broader campus context, a guest speaker from the university's Ad Hoc Committee for AI in the Classroom shared findings from a university-wide survey of 200 students, including data on faculty adoption and dissent. The student leaders then facilitated a discussion among workshop participants (See Figure 3).

Policy Drafting. Student leaders then introduced participants to the policy drafting activity. Following Sanders and Stappers'

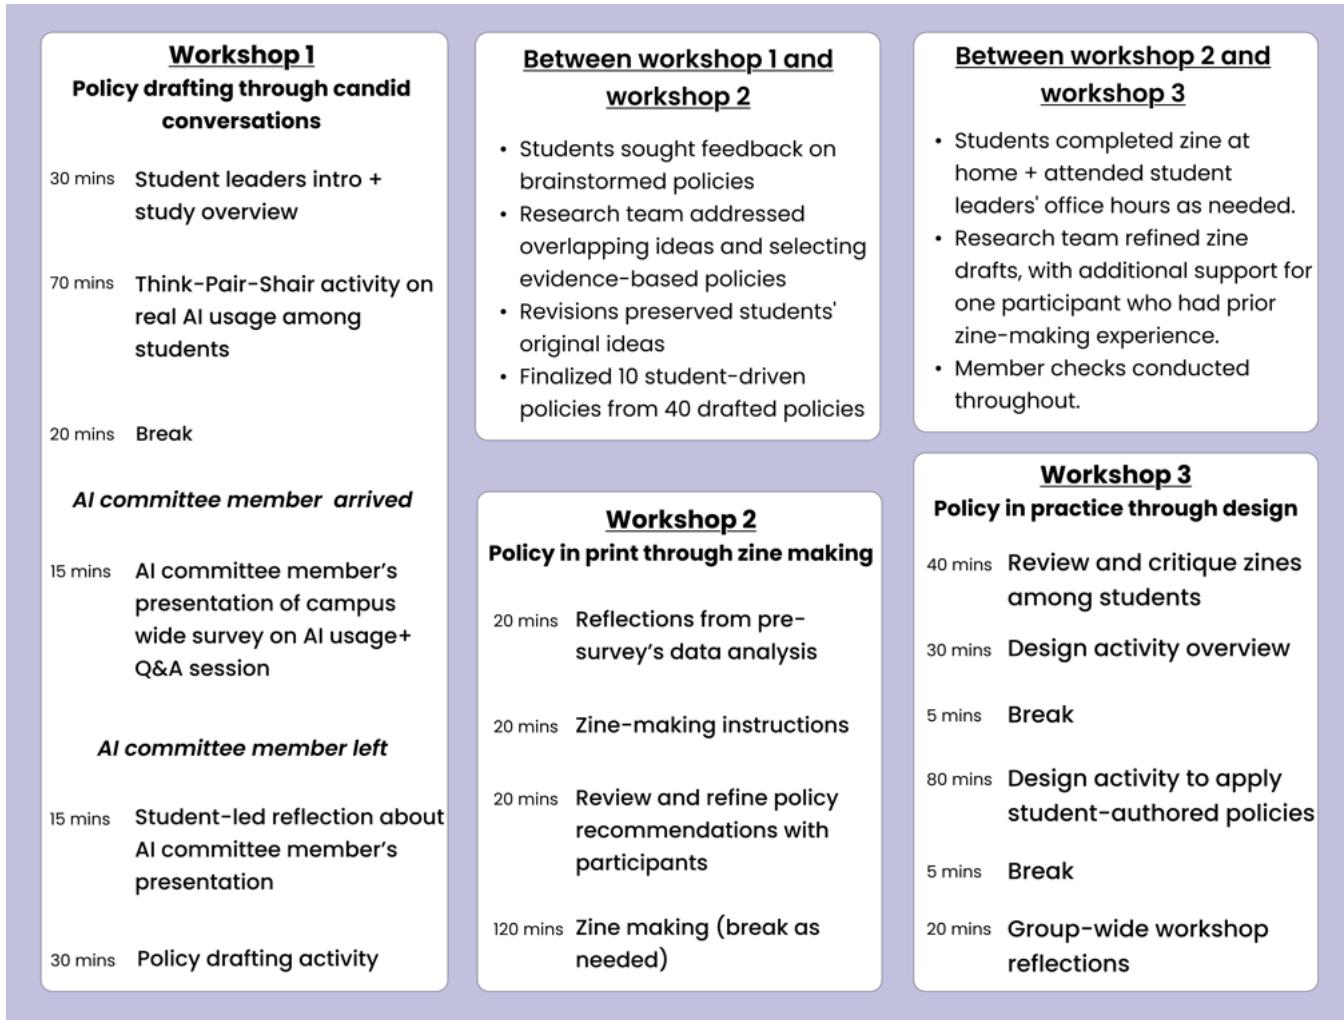


Figure 3: Overview of the three-part workshop series.

recommendations for researcher-as-facilitator in participatory design [108], we treated AI policy design as a domain in which students needed structured entry points rather than a blank slate (based on participants' responses from the pre-workshop survey). As a starting point, student leaders presented 24 potential policy topics, drawn from HCC 629's learning objectives (divergent thinking, feedback), related AI education literature (e.g., agency and equity [65], critical thinking [110], ownership, group work, and social support [24]), ongoing AI policy conversations at UMBC (e.g., AI literacy, grading and rubrics, academic integrity), and the student leaders' own experiences (e.g., perceived double standards in faculty use, English learner considerations). Participants were invited to add additional topics (See Supplemental Materials, Workshop 1 Slides, Slide #34), and one participant did (i.e., Policy #1 "Instructions"). This approach allowed us to seed the workshop with topics that were both pedagogically grounded and locally salient, giving novice policy writers clear, student-relevant entry points while

treating the list as a sensitizing scaffold rather than an exhaustive or prescriptive agenda.

Each participant then selected five topics to draft policy recommendations for. Participants were provided with structured prompts on an activity worksheet that guided both brainstorming and policy formulation. Each policy topic—such as "Divergent Thinking" or "Hypocrisy in Faculty Use"—was presented with a guiding question (e.g., "How can generative AI encourage you to explore diverse ideas and inspire creative approaches to problem-solving?"). Prompting questions helped translate abstract concepts into student-facing questions in everyday language that they could more easily understand how to map their experiences and make policy topics connect to their lived experience. For each theme, worksheets led students through a process to identify a core challenge or tension, describe a relevant learning outcome or classroom scenario, and finally draft a concise policy statement using conditional phrasing and action verbs. This policy drafting activity resulted in 40 initial recommendations (8 participants, 5 prompts each).

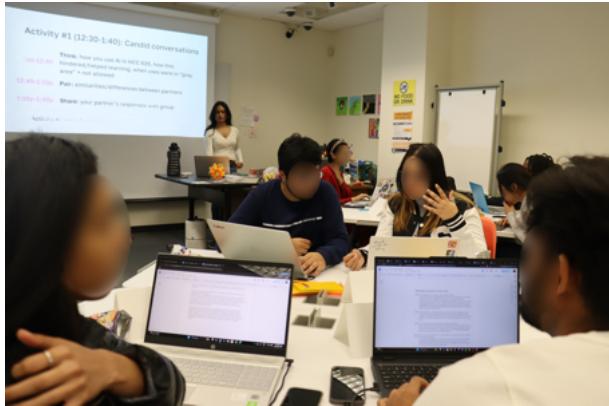


Figure 4: Workshop 1 focused on policy drafting through candid conversations of AI use. Students exchanged reflections on their use of AI tools in a design course the preceding semester.

Between Workshop 1 and Workshop 2. The research team convened three times to refine the 40 initial policy recommendations. While the core sentiment of the student recommendations was preserved, the research team rephrased statements to consolidate vague rhetoric, such as “do not use AI to complete your work.” Widely repeated suggestions, such as many policies focused on academic integrity, were consolidated (See Table 2, Policy #6 “Citing AI Use (or Not)”). The research team contacted participants with clarifications to preserve student intent. The authors reviewed the initial recommendations with related literature to ensure policies were backed by evidence, as well as revising for clarity and novelty. For instance, a participant recommended that “students should only use genAI at later stages of the design process to ensure ideas are their own.” However, generative AI shows potential for supporting early ideation [60]. This synthesis process ultimately resulted in a curated set of 10 student-driven policy recommendations (See Table 2 for this list of final policy recommendations).



Figure 5: Workshop 2 focused on zine-making. Students flip through several exemplar zines while learning about the historical context of zine-making practices.

3.3.2 Workshop 2: Policy in Print through Zine Making. After the research team and participants went through the consolidated and refined policy recommendations, participants started the zine-making process by choosing two of the policies to visualize.

Introduction to Zine Making. Given that most participants had never made a zine before, the research team took steps to onboard students to the zine-making process [54]. Student leaders welcomed a visual arts professor at the university who leverages *zines as research* [53] through their extensive experience of zine making (See Workshop 2 Slides 10–27 in Supplementary Materials). The professor reviewed zine history [99, 102], then facilitated discussion on the relevance of zine making in this project: “Why not just write a typical policy document outlining findings?”; “Why take the time to create a visual representation of a policy?”; “What legacy of zine making is relevant, given the student-driven nature of this workshop series?” Students were challenged to think about what it means to present policies in the form of a zine, rather than typical policy formatting. The research team accommodated multiple approaches to creating zine pages based on preferences and strengths, to further accommodate participants’ familiarity with zine making. Using provided templates, each participant created one digital page in Figma [41] and one physical page on paper. Students were provided materials such as markers, tape, scrapbook and magazine cut-outs, scissors, and glue sticks. Digital tools such as an ASCII art generator [97], the Iconify Figma plugin [5], and generative AI tools such as DALL-E [4] were also provided.

Iterations and Refinements of Initial Zines. While participants had two hours during Workshop 2 to begin zine making, most finished their pages the subsequent week. Participants were allowed to borrow any physical materials provided during the workshop to complete their analog pages. To support participants, student leads held office hours to provide feedback on creative choices and clarify instructions from Workshop 2. This in-between counseling occurred alongside refinement of the draft zines. The research team met five times to review drafts (See Supplemental Materials for the first draft of zines). Member checks were also conducted to garner students’ feedback before sending to print [30, Ch.4].

3.3.3 Workshop 3: Policy in Practice through a Design Activity. In the final workshop, students applied their policies to complete a design activity that resembled their assignment structure from the design studio course. Participants were asked to apply their policy recommendations to redesign generative AI interfaces to better support the design studio course’s learning objectives. Students selected from six generative AI tools spanning image/UI-based (V0 [8], UX Pilot [7], Ulzard [6]) and text-based (ChatGPT [2], Claude [3], Grammarly [50]) categories, ensuring coverage across all tools. Participants completed their redesign on Figma [41], with as many screens as they desired, aiming for higher fidelity redesigns.

Design activity to apply policies. To support their design activity, participants choose from a list of HCC 629’s objectives to optimize their design for as well as a conceptual metaphor to provide additional structure: AI as a co-agent [110], AI as a sensor not a solution [110], AI as Feedback Generator, Personal Tutor, Learner, and Team Coach [88] (See Supplemental Materials for full list of conceptual metaphors). Participants then redesigned the interface of [GenAI tool] by using the conceptual model of [Conceptual Model]

to strengthen [HCC 629's learning objective]. To support this reflection on policies in practice, participants completed a reflection worksheet, which included opportunities to revise policies.

Expert evaluation of interfaces. The participants' redesigned interfaces (included in Supplemental Materials) were then assessed by five senior human-computer interaction faculty at the university, all of whom had taught the design studio course previously and were familiar with the course's learning objectives. Interfaces were evaluated on a Likert scale of 1–3 (1-needs improvement, 2-adequate, 3-strong) with four criteria: learning objective, conceptual model, usability, and overall design quality. Interfaces were ranked by average score; top interfaces may be built in future offerings of HCC 629.

Publishing the zine. The refining stages to produce a camera-ready zine—included in the Appendix (See Appendix A)—extended well beyond the third workshop, until July 2025, when the zines were printed at a student-run printing service on the university's campus. All students involved were invited to celebrate the completion of the zine at a “zine party” in October 2025, where they received copies and discussed the project and ricocheting campus conversations over a provided lunch. The last author implemented these student-driven policies in the proceeding offering of the same course, providing an opportunity to evaluate downstream effects in a live classroom.

Open-sourcing workshop materials. Workshop materials are open-sourced so that other faculty and students can run these activities in different departments (across UMBC but also in other institutions)—see Supplemental Materials.

3.4 Post-Workshop Interviews

After the workshop series, the authors conducted three follow-up interviews with workshop participants to understand how their thinking on the topics discussed had developed since the workshop took place. We also conducted two additional interviews with students who were enrolled in the design studio course but opted not to participate in the workshop series due to dissenting views on AI in order to investigate non-adopters' perspectives on generative AI [138]. We followed a semi-structured interview protocol, with differing questions for workshop participants and non-workshop participants. Workshop participants were asked questions such as: *How has your use of AI in coursework stayed the same or changed since participating in this workshop series?*, and *Have you had any conversations about the workshops, zine, policies, etc. with your peers or instructors?*. Those who did not participate in the workshops were asked questions such as: *What were your concerns with participation in the workshop?* and *How can we ensure students with dissenting views of AI are heard at the policy level?* Interviews were conducted in person or online via WebEx for one hour and interview participants were compensated \$20 (on top of the \$220 compensation for those who also participated in workshops). We acquired informed consent from interview participants, and all data collected from the interviews were securely stored following the university's institutional review board's protocol for protecting research data.

3.5 Data Analysis

Audio recordings from workshops and interviews were transcribed with Otter.AI [95] and manually checked for errors. The joint first authors corrected transcripts and de-identified all materials before sharing them with the faculty PI.

We conducted thematic analysis [25, 26], where all three authors participated in analysis. All authors collaboratively coded the first workshop transcript and created a living codebook. Then, the two joint first authors led the coding of all remaining transcripts (both interviews and workshops), and the faculty PI joined for codebook development, theme refinement, and interpretation of de-identified excerpts. Workshops 2 and 3 were coded independently by the two joint first authors, followed by collaborative sessions where all three authors compared interpretations, discussed discrepancies, and revised the codebook. Calibration meetings occurred weekly during data collection and continued for three months afterward. Disagreements were resolved through constant comparison and privileging participants' phrasing. Our inductive coding approach enabled us to stay close to participants' language and actions; we generated 1,051 initial codes across workshops and interviews. We did not compute inter-rater reliability, consistent with approaches emphasizing the interpretive nature of qualitative coding [83].

Next, we iteratively developed 23 categories capturing recurring patterns (e.g., refusal of AI, concerns about AI use, policy-based concerns) and subcategories (e.g., gray areas, citing requirements, constraints on creativity). To surface relationships, the research team constructed an affinity diagram in Miro [87], with cluster boundaries negotiated in group meetings. Throughout, we maintained an analytic memo corpus (19,408 words) reviewed in weekly meetings to stabilize interpretations. We judged categories conceptually sufficient within this study's bounded context—a single graduate design studio course—when later data elaborated existing categories rather than generating new ones; we treat our analysis as offering a coherent account of this cohort's experiences rather than claiming generalizability to all students or institutions. Finally, we remained reflexive about positionality. The two student co-authors contributed insider perspectives as peers, while the faculty PI remained distanced from identifiable transcripts to reduce the influence of faculty authority. This insider–outsider pairing supported proximity to student voice while preserving critical distance.

4 Findings

We structure this section around three key findings: 1) how the workshops supported candid disclosure of AI practices and rationale, grounded policy discussions in recent lived experiences, and scaffolded policy articulation resulting in, 2) ten student-authored policy recommendations, and finally, 3) how both the policies and process of participation sparked reflection among participants—and the campus community more broadly—towards more intentional AI use and the importance of participatory governance of AI.

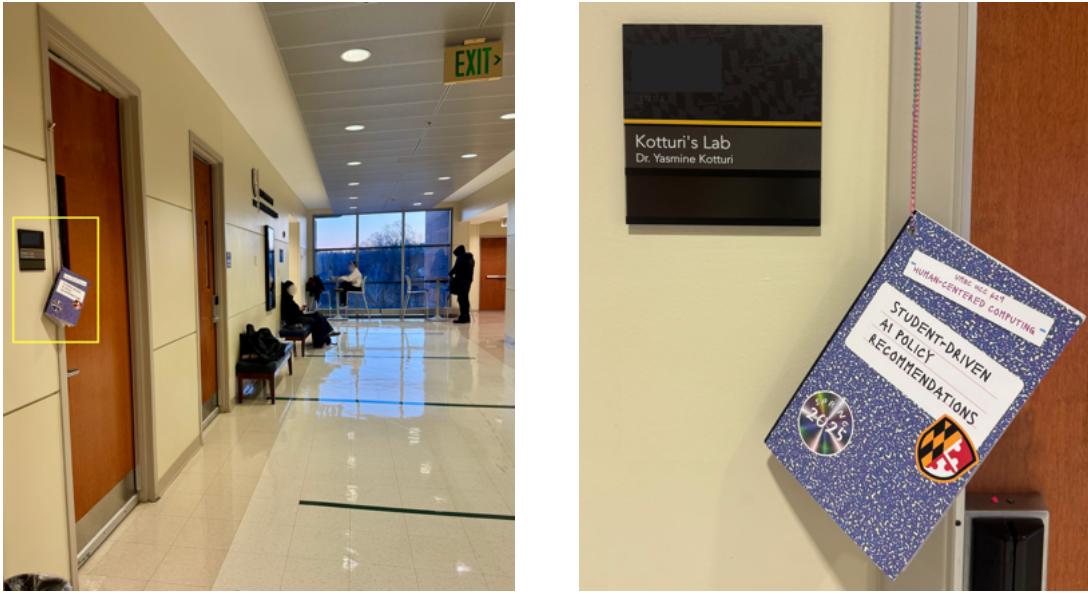


Figure 6: The Student-Driven AI Policy Recommendation zine displayed outside of the research team's lab (right), in the main hallway of the building (left). This central location facilitated student discussions beyond the bounds of the workshops. In addition, over 300 copies of the zine were printed and dispersed throughout UMBC's campus.

4.1 Setting the Stage for Policy Making: Candid Disclosure of AI Use among Students

The workshops created conditions for candid peer dialogue that participants described as uncommon in their everyday academic experiences (P1, P2, P6, P8). As P6 noted:

“It’s not common to talk about what we’re all using and doing with AI. This is a nice opportunity to hear those perspectives.”

P8 emphasized that despite AI’s controversial status, the workshop provided space to “*discuss openly and freely*”: “*I think AI is controversial for some part, but it’s very essential these days, many companies and many students and many people are using it. So I think it is a great opportunity to discuss openly and freely.*”

Notably, participants were often unaware of how their peers used AI until the workshops surfaced these practices. When a member from the university’s AI committee presented campus-wide survey data, participants expressed surprise at the breadth of adoption. P1 shared, “*I find it surprising that even visual arts students are using genAI quite a lot. I didn’t really expect that.*” P2 followed up with “*The psychology students are using the GenAI. That’s a good thing and a bad thing.*”—a comment that sparked reflection on AI adoption in fields like psychology, where students might go on to become therapists; for P2, this realization was unsettling.

While faculty concerns regarding student overreliance are well-charted in the related literature, thick descriptions of students’ concerns towards overreliance of AI systems are less understood. Most participants described their use to overrelying on AI and expressed concerns about this dependency [P1, P2, P5-P8]. At the crux of this overreliance was the pressure of deadlines and feeling

overextended. Due to the time pressure or their self-prescribed laziness, they ended up using AI. Sometimes, participants considered AI as a safety net, leading them to habitual cognitive offloading and hindering independent idea and content generation. For instance, P7 shared: “*We are busy in those cases wherein we need content ideas, brainstorming. Our mind has gone into lazy mode and it has become a habit.*” P8, who used AI for reading summarization, expressed concern because they relied on AI to read papers rather than their own reading comprehension; they felt guilty delegating work to AI that instructors expected them to do. Participants shared various concerns about using AI as it related to their mastery of the course’s design material. For instance, regarding their comprehension of concepts, their concerns echoed the faculty and administrators who typically dictate policy: they warned that certain uses—generating responses entirely [P1, P6], direct answering [P3], unexamined edits [P2], and summarization [P8]—displace the thinking that leads to conceptual grasp.

4.1.1 Hidden Practices: Policy Violations and Gray Areas. There were instances where participants reflected on how and why they used AI in ways that they knew clearly violated their course’s AI policy [P1, P3, P5, P7]. Some described relying almost entirely on ChatGPT (e.g., “80-90%”, P7) to complete assignments with little personal contribution, while others shared how they incorporated AI-generated outputs without citation despite clear prohibitions. Participants concealed their usage because they either knew it was unauthorized, were unsure if it was allowed, or felt guilty about being unable to complete the assignment independently. One participant noted that even in an assignment designed to critique AI’s perspective, classmates turned to ChatGPT to write their critiques

[P3]. P5 described a specific instance where they used AI in a way that was disallowed:

"I used [AI] to write a discussion [post]. And that day...I just wasn't feeling it, and it didn't involve any sources at all. I just said [to the AI tool], 'Answer this for me.' I just pasted it. And I know some professors have AI detectors, so they probably just put it through there. I was like, 'Yeah, you know what, I'll take the L that day.' So that's what happened...they just put in the comments, 'I could just tell this wasn't you...' Instead of 100[%], they gave me an 80[%]."

While this incident did not occur in HCC 629, it points to an interesting phenomenon: P5 received an 80% on an assignment that was entirely AI-generated and in clear violation of their course's policy. This mismatch between the rules and application of penalties may be contributing to the confusion and disarray, and warrants specific attention, as we further explore in the discussion section.

4.2 Ten Student-Driven AI Policy Recommendations

Facilitating candid conversations allowed participants to brainstorm, apply, and refine policy recommendations. In this section, we define each of the ten policies and provide details such as disagreements among participants or challenges experienced when participants put their policies into practice.

4.2.1 Policy Recommendation #1 Instructions: *Instructors should include guidelines with concrete examples of acceptable and not acceptable use of AI for each assignment.* See Figure A1 in the Appendix for corresponding zine page.

Participants felt frustrated by the lack of clear guidelines on acceptable AI use. The AI policy in HCC 629 allowed students "to use AI lightly" but required citation via chat logs with "strict enforcement", mirroring a common vagueness found in existing AI policies [114, 123]. Students were uncertain what this meant in practice. After re-reading this policy, P3 still wondered: "*how much can we use AI? What does 'light use' mean?*" As P1 and P2 emphasized, sometimes their AI tool would alter or generate too much, taking their use into a gray area without their approval, making it hard to "undo."

This ambiguity worsened when enforcement did not align with the stated policy. As noted earlier, P5 received 80% for a submission entirely generated by AI in clear violation of course policy. Similarly, P10 reflected on peers who used AI to create video submissions despite explicit prohibitions, yet faced no consequences: "...*So what [did] we learn, that [students] can use AI in the video, even though they were told not to use it, and they can still pass the class. Next semester, if someone says, 'Hey, don't use AI for this assignment, otherwise you won't pass the class.' And if they still use AI for the assignment, they're expecting to pass that class.*" This disconnect between policy and enforcement made AI guidelines even more tenuous.

To address this ambiguity, P1 suggested providing concrete examples:

"[Instructors] can provide acceptable and unacceptable use of AI...in the starting of the class. Run through

a presentation and give [students] acceptable and unacceptable AI examples."

P1 noted this could also strengthen esprit de corps: "*I feel like it should be done because it makes life easier. If you leave students on their own, they probably will be confused.*" P6 agreed that examples would reduce fear around citation: "*If someone is using a lot of AI and they fear not citing it, in order to remove that fear, like [P1] said, a brief example would do it.*"

Finally, participants noted that blanket course-level policies often lacked the granularity needed across assignment types. Tool-specific policies—such as banning ChatGPT but allowing Grammarly—were also becoming outdated as more platforms quietly integrate generative AI features, reinforcing the need for assignment-level guidance.

4.2.2 Policy Recommendation #2 Ownership: *We should only use AI for up to 50% of our work on any given assignment, so that the majority reflects our own ideas and effort.* See Figure A2 in the Appendix for corresponding zine page.

Ownership—also referred to as "authorship" and "accountability" throughout the workshops—generated rich, sometimes contentious discussion among participants. Students naturally framed ownership in terms of percentage: how much could be AI-generated while still being considered their work? There was disagreement on this threshold, captured in the corresponding zine page on "post-it notes". P4 argued that at most 30% should be AI-generated; P1, P6, and P9 proposed 50%, concerned that higher thresholds would be exploited; P7 and P3 argued the percentage was irrelevant as long as the original ideas came from the student. Ultimately, 50% emerged as a consensus halfway point.

However, after applying this policy in the final workshop activity, P3 shared how this policy was too rigid in their post-workshop reflection:

"Honestly, the biggest issue is that [the ownership policy] feels a bit rigid, especially when you're not using AI to do the thinking for you, but just to move faster. Like in my case, I had a clear mental model, and the AI tool just helped me skip the dragging-and-dropping part. But if we follow the current rules strictly, that might still count as 'more than half AI-generated,' even though the actual decisions came from me. That's where it gets tricky."

P3 disambiguated different kinds of use, noting that not all AI use should be policed the same way—distinguishing use that helps "move faster" when one already has a clear vision from use where AI acts as a collaborator (a concept described in detail in [110]): "...*I wasn't blindly accepting what AI gave me. I already had the structure in mind, and AI just helped speed it up. It felt like streamlining, not outsourcing.*"

P1 noted that as heavy AI use becomes normalized in industry, such thresholds may need reevaluation to support employability. P6 pointed to the labor required to maintain ownership when AI suggestions were strong: "*I had to consciously revise and rewrite AI-generated text to make sure it reflected my own thinking. Balancing usefulness with ownership required effort, especially when the AI*

suggestions were strong.” While participants were open to following this policy, they admitted compliance would depend on context.

Finally, another challenge was measurement: how can someone calculate the percentage of AI use across multiple phases of the design process? Relatedly, and connecting this issue to Policy #1 on Instructions, P4 noted that examples would clarify what 50% actually looks like in practice: *“This policy needs to be explained to all students with this image generation and text generation example to show what less than 50% actually means, otherwise, it didn’t even make sense to me.”* So while all participants agreed that ownership over their own work was essential, the particulars of how to conceptualize and quantify this proved challenging; we revisit and delve into this further in the discussion section.

4.2.3 Policy Recommendation #3 Divergent Thinking: *When using AI for brainstorming, we should push ourselves to explore alternatives, surprising directions, or ideas that feel more personal and meaningful to us.* See Figure A3 in the Appendix for corresponding zine page.

Participants gravitated toward divergent thinking—the essential stage in design that focuses on generating many ideas before converging on a subset to pursue [122]. Participants noted that AI use could both hinder self-expression and creativity [P1, P3, P4, P5, P7, P8] and support it [P2, P6], depending on context—especially when deadlines or workload would otherwise block deeper engagement.

When applying this policy during the design activity, several participants reflected on using AI to support ideation. P1, P2, and P5 noted that the policy encouraged them to prompt AI to branch ideas, explore alternatives, and resist shallow responses. P2 shared:

“I used GenAI primarily as a creative amplifier, especially in the ideation phase...My approach remained iterative. I would prompt the model, reflect on its output, and then reframe or refine ideas to suit my intended direction. The back-and-forth helped me stretch my thinking without losing authorship over the outcome.”

P8 felt AI could theoretically help when peers are unavailable for feedback, but found the outputs too generic in practice: *“If I don’t have any peers to brainstorm together, [then] I could ask [AI]. [But] I feel like it gave me too broad, too general viewpoints.”*

P10 offered a dissenting perspective, questioning whether AI-assisted design undermines skill development:

“As long as you have the fundamentals down—you should at least have those down—if you’re using AI for every step of the design process, are you doing the work? Because you can say that, ‘based on my personal taste, I like to do this,’ but then are you doing the work? Or is AI doing the work at that point?”

P10’s concern points to a consistent challenge of overreliance on AI and lack of skill building: AI can prevent students from developing the foundational skills they need to do the work independently.

4.2.4 Policy Recommendation #4 Job Skills: *We should be provided opportunities to learn how to use AI in ways that reflect how AI is used in real workplaces through coursework. Instructors should stay updated on how AI is used (and regulated) in industry.* See Figure A4 in the Appendix for corresponding zine page.

Participants noted that most students learned to use AI on their own rather than through formal instruction, resulting in significant knowledge gaps: P7 was concerned about identifying AI-generated bias; P1 felt unprepared to write effective prompts; P3 noted that lack of training contributed to academic integrity concerns. Currently at UMBC, as at many institutions, AI training is not provided to students.

Participants argued that any training should center employability. P1 shared how their roommate’s new employer expected all employees to use AI from “day one”. P6 reflected that AI skills were now required for UX job applications—the career path most participants intended to pursue. P7 added that AI literacy was also important for research careers and Ph.D. programs. Alongside these needs, participants grappled with the lack of support their university and instructors provided.

However, P5 raised an important caveat: in some jobs, such as government positions, AI use is strictly prohibited. Given the university’s geographic proximity to government employers, this constraint was relevant for many alumni: *“...what’s the point of teaching it if we can’t even use it in the first place?”* P5 emphasized that any training or guidance should account for these job-specific considerations.

P10 pushed back on this policy entirely, arguing that the course was meant to teach fundamental design skills, which should not be supplemented with AI assistance; students should be able to complete tasks independently in case AI becomes unavailable—such as when transitioning to a job where AI use is prohibited, as P5 described.

4.2.5 Policy Recommendation #5 Bias: *We should check if AI-generated ideas include any stereotypes or biased assumptions, such as by asking if any perspective or voice is missing in the response received.* See Figure A5 in the Appendix for corresponding zine page.

“Bias” was one of the policy topics participants did not initially gravitate toward—not due to lack of interest, but lack of practical support. When reflecting on this policy in practice, P1 struggled with the labor required, noting no commercial tools exist to help: *“I had to carefully go through the content, analyze language and representation, and possibly cross-check against standards or guidelines of the work to ensure that no biased assumptions or stereotypes were present, which was time-consuming.”* Other participants found the policy easier to follow. P2 shared: *“I stayed away from anything that felt biased or vague—if something felt off or too generic, I just changed it.”* P4 took a different approach, relying on AI itself to check for bias: *“I always ask AI if there is any bias and it is usually quick to detect that, so that’s why I like adhering to this policy.”* While participants differed in their strategies for addressing bias, their reflections point to a broader need for accessible tools that help students identify and critically assess biased content in AI outputs.

4.2.6 Policy Recommendation #6 Citing AI Use (or Not): *It is important for us to credit AI where credit is due. However, sharing chat logs, as many instructors currently require, is tedious and ineffective for both students and graders. Instead, we should share a 2-3 sentence summary for each submission describing how and why we used AI.* See Figure A6 in the Appendix for corresponding zine page.

Citing AI use was one of the most debated topics across workshops and post-workshop reflections. Almost all participants recognized that citing via chat logs was tedious; P5 called it “a big ask,” and P7 believed AI use should not require citation at all.

P9 offered a counterpoint, arguing that chat logs—even partial ones—would make dishonesty harder and give graders a clearer sense of how much AI was used. While acknowledging that instructors may not read every log in detail, P9 felt that even skimming would reveal red flags and considered logs more transparent than brief written summaries.

P6 highlighted a potential benefit of citation: if students cited their AI tools, others could learn what tools and approaches their peers were using. As P6 noted: *“I think if they get a hold of it, or if they see someone else using something different from them. That’s where they realize, if they cite it, it’s gonna be useful for them and others too, when they collaborate and when they realize the benefits of it.”*

However, most participants resisted citation requirements, pointing to disconnects between policy and intention. P4 shared how they circumvented the requirement by prompting AI to cite itself:

“One of the policies is basically, ‘[cite] everything everywhere we use it.’ Basically, if you tell GenAI ‘Consider the policy and then cite yourself appropriately wherever you are.’ It does that whether it’s citing or paraphrasing or putting the references that will cite itself.”

P5 and P7 raised a deeper concern: citing AI use meant, for them, revealing how little of the assignment was completed independently. P5 asked:

“Is there such a thing as being too transparent? Because if the professor is asking [you] to cite the use of AI, and you have AI do the whole project for you, and you cite that, then you’re cooked...But versus, if you use AI to edit a line of your paragraph, or a line of code, and you still cite it, you get points taken off.”

Tucked behind what appears to be overreliance on AI, P5 is also pointing to the uneven enforcement of AI restrictions, within and across their courses. Given this unevenness, it further pushed them towards more opaque use of AI (the logic being: if I am going to get dinged for any AI use, big or small, then there is no incentive to report truthfully. P7 candidly chimed in on a related note:

“So one major problem I had in HCC 629[was that] sometimes I used 80 to 90% of ChatGPT to do the assignments, [and] only my [completed portion was] 10%. So I had a fear, [if I have to cite], I cannot cite everything, right? It will look like I’m just using ChatGPT for the assignment. It shows that you have done nothing, right?”

Both P5 and P7’s candid reflections highlight a key motivation for understanding why students do not include chat logs and report AI usage, even when more experimental policies are in place. Given this ineffective citation mechanism, P1, P2, and P6 proposed an alternative solution which may prompt more reflection and critical use, while also making adherence to citation policies more likely: include a short summary describing how and why AI was used. To

capture the range of views, the corresponding zine page included post-it notes with dissenting perspectives (see Figure A6 in the Appendix, page 25).

4.2.7 **Policy Recommendation #7 Hypocrisy in Faculty Use:** *If instructors expect us to be transparent about our use of generative AI, we expect the same transparency from them when it comes to how they use AI in teaching, grading, or creating assignments.* See Figure A7 in the Appendix for corresponding zine page.

Students perceived a double standard: their AI use is highly regulated, while those who dictate policy face no similar requirements. P4 argued that transparency should apply equally to students and faculty, adhering to the golden rule. Beyond equitable treatment, students wanted to understand how instructors’ approaches might be changing in response to AI—and what that meant for their learning. One student leader reflected that if faculty use AI to free up time on rote tasks, that time should be reinvested in one-on-one instruction.

P9 took a stronger stance, arguing that faculty should demonstrate expertise by producing teaching materials themselves rather than “cheating” with AI:

“If you’re using [AI] to teach, get your Ph.D. revoked, honestly. You should know how to teach the content that you wrote a dissertation on. I don’t think that would be in any way acceptable.”

While P9 was fine with minor uses like grammar or formatting, they emphasized that students pay tuition for professors’ knowledge, not AI’s output.

4.2.8 **Policy Recommendation #8 English Learners:** *For those of us who are English learners, we should be encouraged to use AI to support our English proficiency in writing by asking for refinement of our text. Also, we should ask for an explanation of the refinements made to further our language proficiency.* See Figure A8 in the Appendix for corresponding zine page.

P1, P2, P7, and P8 expressed difficulties conveying their thoughts in English, given that it was not their primary language. With limited vocabulary and grammatical structures, their writing often felt overly simple, and they expressed frustration at being unable to articulate ideas as clearly as they wished. For these participants, AI tools were almost necessary to support writing at an academic level.

Participants emphasized that policies should explicitly recognize AI’s role in helping English learners understand context and fully express ideas—including emotional nuance—without penalty. P2 noted that the level of fluency should be factored into any policy, a point echoed by one of the student leads, also an English learner. P5 and P8 observed that even native English speakers use AI for grammar and spelling support, suggesting the line between acceptable and unacceptable use is blurry for everyone.

However, P9 cautioned against treating AI as a language learning tool. They acknowledged that AI could help when students struggle to put thoughts into English, but warned that English learners might accept AI suggestions at face value—even when the wording or tone is inaccurate—inadvertently learning mistakes. As P9 put it: *“You don’t know the words to say, and AI is giving you something. How do you know AI is giving you the actual words you want to say?”*

Ultimately, this policy aimed to ensure that students less proficient in English are not penalized for leveraging AI, while recognizing the need for critical evaluation of AI suggestions.

4.2.9 Policy Recommendation #9 Feedback: *We can use AI to help us better understand and respond to peers' feedback, especially when revising our work.* See Figure A9 in the Appendix for corresponding zine page.

This recommendation received unanimous agreement during the initial policy review. By reframing AI as a feedback intermediary, the policy positioned AI not as a replacement for peer insight, but as an aid in making feedback more digestible and growth-oriented—especially in classrooms where social dynamics inhibit candid exchange. Because participants did not exchange feedback on each other's design assignments in Workshop 3, they were unable to put this policy into practice, resulting in less discussion to capture. One emergent challenge during the making of the zine was how to visually represent this policy. Students proposed depicting AI as a “shield” for feedback—initially imagined as a Captain America-style shield deflecting harsh comments, then refined into a filter that softens harmful feedback while translating vague critiques into actionable guidance. This metaphor framed AI as a mediator that scaffolds emotionally intelligent exchanges, addressing concerns raised throughout the workshops: discomfort with critique, unclear peer review norms, and the emotional labor of giving and receiving feedback.

4.2.10 Policy Recommendation #10 Equity: *Everyone in class should have access to the same AI tools or models for each assignment to ensure fairness.* See Figure A10 in the Appendix for corresponding zine page.

This policy addressed potential inequities in access, as some students can afford more powerful AI tiers than others. Interestingly, the original version of this policy was more prescriptive: “the same AI tool and model should be used by all students for each assignment.” However, upon applying the policies in the third workshop, participants noted loopholes and found it too restrictive. P4 shared: if a policy mandates identical tools, students who build their own AI tools—demonstrating technical creativity—would technically be in violation. P7 emphasized that standardizing access—providing a baseline set of tools to all students—would ensure fairness and help students navigate the rapidly expanding AI landscape. However, P3 noted the challenge of maintaining such a list: “*There are a lot of AI tools right now. People don't even know the existence of them.*” P10 questioned whether students would receive institutional AI subscriptions and what ethical standards should guide that selection. P9, who does not use AI tools, argued it would be unfair to pay for a subscription they would not use—similar to the gym fee the university already charges them.

4.3 From Policy to Practice: Shifts Toward Intentional Use

Participants shared that the zine-making process created a sense of permission to question and debate classroom AI policies. One of the most notable shifts after the workshop series was participants' increased intentionality in how they used AI. Many described changes in their practices, often referencing the conceptual metaphors from

the design activity—such as AI could be used as a reflective or critical design partner [P1, P7], a feedback generator [P1], or a collaborator [P1, P2, P6].

Strategies for intentional use included shifting from passive to reflective collaboration [P1], bringing more critical insight to their work [P2], and prompting AI iteratively throughout a task rather than with a single prompt [P6]. P2 shared: “*This time, I used [AI] more intentionally and strategically. I treated GenAI as a collaborative thought partner to question my assumptions, reframe ideas, and push past obvious solutions.*” P6 echoed this: “*GenAI served as a thinking partner helping to speed up decision-making without doing the thinking for me.*”

Rather than using AI simply to complete tasks, participants worked with it collaboratively—approaching use with intentionality. P7 suggested that institution-wide AI policies should serve not just to penalize misuse, but to increase awareness and intentionality:

“I learned it is very important to have AI policies in the institute, because students are not aware how they're using AI, how much they want to use, and how they can use Gen AI [tools]. Having such policies and having such guidance and instructions will help students to be aware of the future opportunities, future challenges and risks.”

Involving students in the process of creating policy allowed them to reimagine how AI could support—rather than undermine—their learning, a point which we revisit in the discussion section.

4.3.1 Ripple Effects across the Campus. The zine, policy recommendations, and ethos of centering student perspectives circulated beyond the workshop context, catalyzing continued conversation among faculty, students, and senior administrators. Faculty across disciplines drew on the zine as a prompt for classroom practice. Two HCI faculty guided discussions on AI policies with their students; a computer science professor brought the zine as a starting point for co-designing AI policy during the first week of an AI ethics course. Other faculty encountered the zine hanging outside the research team's lab (See Figure 6). The project also seeded a new research collaborations: one participant began working with faculty on a related project, and one student leader is now collaborating with the university library to run a campus-wide series of these workshops across disciplines, extending the participatory model beyond a single domain and course. The zine has since been archived in UMBC's Library—both physical and digital versions²—ensuring that student perspectives persist in institutional memory beyond the workshop series.

However, these ripple effects also revealed persistent misunderstandings of the project's goals. When the project was presented to two faculty members in the research team's department, they initially interpreted the zine as a mechanism for increasing compliance—helping students adhere to rules and improve grades. We stress that the goal of this project is not strictly enforcement-oriented but rather centering students as experts, or “lead users,” within this context [130]. These misreadings underscore how deeply ingrained top-down framings of AI policy remain, even among faculty sympathetic to student perspectives. By tracing these follow-on

²<https://lib.guides.umbc.edu/c.php?g=1475372&p=11187429>

effects, we see how the zine extended workshop conversations into everyday pedagogical practice—while also surfacing the ongoing work required to shift institutional norms.

5 Discussion

In this discussion, we re-situate the student-driven policy recommendations in contemporary discourse around university AI policy—a rapidly evolving space [68, 101]. In particular, we consider how these student recommendations fit into the existing spectrum from experimental policies—where AI is integrated into classrooms as a supportive tool—to restrictive policies—where AI is viewed primarily as a threat to traditional learning methods [119]. To do so, we first unpack what practical challenges exist when implementing students' policies (such as when policies might not be synergistic with learning goals). Then, we consider what the value of student-driven recommendations may be (while remaining cognizant of the aforementioned challenges), such as how the process of participation, rather than solely the resulting policies, is an essential and transferable component of this work. Finally, we conclude with recommendations for how to engage students' perspectives on AI policy in higher education.

5.1 Challenges with Implementing Student-Driven Policies

In this section, we discuss challenges with implementing student-driven policies, beginning with challenges related to feasibility, including enforcement, overhead, and lack of data. Then, we discussed challenges that arise when there may be an incongruence between students' policies and course learning goals.

5.1.1 Feasibility challenges with student-driven policies. We consider several of the students' policies—“Equity”, “Ownership”, “Job Skills”, and “Instructions”—which may not be feasible to implement for various reasons such as enforcement difficulties, technical limitations, large overhead for instructors, or simply, lack of know-how. To begin, with Policy #10 on “Equity” (See Section 4.2.10), participants sought equitable tool access across their peers, so that no one student had a particular advantage by accessing expensive or lesser-known AI tools. In some ways, restrictive AI policies address this issue: by disallowing all use of AI, students do not have uneven access to AI tools [35]. Instructors with experimental policies, on the other hand, have attempted to tackle this issue by co-creating statements with students at the beginning of their course, where all students agree upon which tools they will and will not use [121].

Ultimately, one of the practical challenges across all of these approaches towards equity is enforcement [48]: how to make sure students do not use the tools that their policies, either co-created or not, disallow them from using. In order to understand how to address issues of enforcement, we argue it is essential to dig into the reasons why students do not abide by policies in the first place. In the case of equity, one participant shared that they created their own AI tool to aid with their assignments, and it felt unfair to not allow them to use this, even though other students could not access it. By understanding the reasons that students have for not adhering to course policy, instructors can decide an appropriate

course of action—in this case, can the student provide an in-class demonstration of the AI tool they built, and even extend use to their peers? Or should the instructor facilitate a discussion for why such personalized tools give that student an unfair advantage? Not only is such transparency likely to support enforcement, it may also foster more critical thinking and awareness around use of AI tools [29, 134].

Policy #2 “Ownership” (See Section 4.2.2) was another policy where students discussed enforcement challenges themselves during workshops: the 50% threshold, for instance, assumes an AI's contribution can be quantified and distinguished from a student's contribution. In this case, enforcement relates to a longstanding challenge connected to remixing in design, often summarized in the phrase: *all design is redesign* [71]. Participants considered how it is not clear where their ideas stop, and the AI's ideas start, especially when rapidly iterating with an AI. Restrictive policies bypass challenges of student ownership by not allowing any use in order to preserve student ownership [127], whereas experimental policies often assume students' responsibility for anything they submit, regardless of what tools have been used [126]. In design pedagogy, understanding the importance of process continues to be a key learning objective [122], and process-tracing tools like BoodleBox [1] may provide one path for exploration in future work when attempting to address this policy's implementation challenge.

Another policy with feasibility challenges was Policy #4 “Job skills” (See Section 4.2.4), as participants wanted instructors to update course materials to include information about what AI skills they need to have in order to land their desired job; often in response to repeatedly witnessing headlines about AI's takeover of knowledge work. On the one hand, this seems like a simple and reasonable ask. Increasingly, universities with experimental policies, such as Ohio State's “AI Fluency” initiative [93], claim connections between their revamped courses, AI literacy, and career readiness. However, it is unlikely to be that simple: currently there exists minimal empirical evidence to support claims around AI literacy and AI skills [104]. In many contexts beyond traditional ML/AI courses, where students learn fundamentals to build, train, and deploy AI models, it is not clear which skills are central and worth pursuing (recall how short-lived the 2023 most in-demand career of prompt engineering was [21]). In design-related careers, for instance, AI's infiltration is leading experts to both critically evaluate AI-generated outputs, which may be unreliable and misaligned with user needs, and develop effective human-AI collaboration skills [73].

Taken together, even as data becomes available on how AI is reshaping skills and careers, and what AI literacy actually entails [104], it is then a non-trivial task for instructors to re-imagine and rebuild curricula to incorporate these learnings, especially given the often-lack of institutional support offered to instructors [16, 82]. Similarly, Policy #1 Instructions (See Section 4.2.1) beckons additional instructor labor. The overhead to implement assignment-specific instructions for AI use, while ensuring academic integrity, and as AI tools' capabilities are changing rapidly, is a herculean task.

5.1.2 What Happens when Student-Driven Policies are not Aligned with Learning Goals? While scholarship on universities' AI policy is a fast-moving and capricious space, there seems to be one clear signal: students tend to be more lenient than faculty on AI use and more supportive of AI technologies writ large [103, 106, 112]. One reason for this difference is that instructors are often concerned about students' over-reliance on AI tools, or that AI use will likely impede student learning by cognitive offloading [64]. In our study, this tension exists in several of the students' policies, such as "Citing AI use (or Not)", "English Learners", and "Ownership".

When it comes to the need to cite the AI tools students use, one relevant learning goal is clear: students need to know how to cite and paraphrase reputable sources [39, 107]. Most of the UMBC courses participants had taken required them to submit all chat logs from AI tools as part of their bibliography (similarly to many existing experimental policies [27, 33]). But students noted the tedium, and even hinted at the irrelevance, of this request: not all AI systems they use are embedded in chat-based tools which afford log sharing (See Policy #6 "Citing AI Use (or Not) in 4.2.6). Instead, students, after lengthy debate, compromised on including a summary of which tools were used and how at the end of the assignment—more or less a process and tools statement [120]. Still, there was strong sentiment that citing AI tools (even in this relatively simple way) is an undesirable and overly tedious task—a stance which is in clear opposition to an essential learning goal of university education.

Policies #8 "English Learners" (See 4.2.8) and #2 "Ownership" (See 4.2.2) followed a similar pattern where students' policies may diverge from learning goals. For instance, some restrictive policies specifically flag that non-native English speakers should not use these AI tools in any way to overcome a language gap [84], whereas other universities with more experimental policies, such as at Teachers College Columbia University, state that these tools may present unique opportunities to support this population [92] (which some new startups hope to capitalize on [70]).

Restrictive policies which touch on ideas of ownership typically disallow use of AI because the very nature of AI tool use complicates students' ability to take responsibility for their work [118]. In our workshops, students' policies supported liberal use of AI among English learners, and complicated notions of ownership to include AI-generated outputs as students' original work—but such ideas may be orthogonal to learning objectives. In these cases, it may behoove faculty not to solicit students' perspectives on AI use in order to safeguard their own learning. But, we argue, rather than disallowing use sans explanation, it is critical to provide an accompanying rationale for such restrictions in order to encourage students to stay within the bounds of acceptable use.

5.2 The Benefits of Student-Driven Policies

In this section, we detail what the benefits of student-driven policies may be, while remaining cognizant of the challenges described above. We argue there are two key beneficial aspects to student-driven policies: 1) the new perspectives students contributed by reflecting on their lived experiences as early adopters of generative AI technologies, and 2) that the process of participation itself is a

valuable step towards a future of AI governance in higher education that serves all stakeholders. For the latter, we provide actionable suggestions for how to scaffold such participation.

5.2.1 Students Contribute New Policy Ideas and New Perspectives on Existing Policies. During workshops, student-led discussions enabled fresh takes on existing policy topics as well as new policy topics that may be important for policy designers to consider. Take Policy #9 on "Feedback" for example (See Section 4.2.9): while some experimental AI policies state how AI may be used by instructors to provide feedback to students on their assignments [128], students in our workshops discussed how they could use AI to digest instructor or peer feedback and translate it into actionable suggestions or to make it less harsh (an issue with peer feedback in particular [57, 66]).

In addition to fresh perspectives on existing policies, both student participants and student leaders—while in conversation with workshop participants—captured policy topics not currently included in university policies. For instance, while both experimental and restrictive policies emphasize university-wide or course-wide AI policies [13], participants emphasized how policies needed to also be *assignment-specific* and provide examples of what acceptable AI use is; captured in Policy #1 Instructions (See Section 4.2.1). Participants sought more granular information about what is and is not allowed to assuage anxieties around academic integrity accusations, especially given how participants pointed out how some existing policies are overly vague [48].

Poignantly, participants discussed at length their frustrations about existing AI policies, especially restrictive policies that ban AI use among students, but where instructors are permitted to use AI [58]. In other words, participants discussed a double standard (what student leaders referred to as "hypocrisy in faculty use [of AI]" (See Section 4.2.7)—yet another policy consideration which is often absent in university AI governance discourse, both restrictive and experimental [58, 67]. Increasing the diversity of stakeholders who design and govern the systems that impact them leads to more effective design and governance outcomes [52, 91]. In our work, we investigated what this could like in AI governance in higher education, where students' perspectives are not included in the decisions that govern them [24, 101]. By engaging student perspectives and empowering an overlooked stakeholder group (albeit within the context of one course at one university) led to new policy ideas and fresh perspectives on existing policies.

5.2.2 Additional Benefits of Student Participation beyond Resulting Policies. Thus far, we have focused on the products of student participation—the policies themselves—and their associated benefits and challenges. We now turn to the process of participation, and argue that engaging students in AI governance carries value beyond the resulting recommendations such as improved student engagement, critical thinking, and compliance, improved student-instructor relationship building (in a moment where trust is fraying bi-directionally [49, 77]) and even the opportunity for impact beyond the bounds of the workshops by shifting students from passive policy subjects to reflexive public's capable of ongoing intervention in institutional AI governance [34].

The “Students-as-Partners” (SaP) paradigm shows that engaging students as active collaborators in teaching and learning practices comes with a host of additional benefits beyond improved efficacy such as enhanced perspective taking and increased academic motivation [32, 55, 80]. This is because co-creation shifts classroom dynamics from hierarchical to collaborative, and calls on students to triangulate and reflect on the status quo, their lived experiences, and imaginatively brainstorm improved futures [32]. In our study, engaging students-as-partners facilitated students’ understanding of the tensions that exist within AI policy making decisions, as well as the reasoning behind existing policies. In addition, engaging students-as-partners facilitated our understanding as researchers about the kinds of subtle barriers that may exist when fostering student participation in AI governance: recall participants’ hesitation to ask clarifying questions about their instructors’ AI policies, as they feared this could be interpreted as an admission of misconduct: *questions are confessions*.

While the scope of our study is limited to one graduate design course at one university, these initial explorations point to an exciting potential of students-as-partners in AI governance in order to facilitate their engagement, critical thinking, and even compliance; when a student understands the reasoning behind a rule, they are much more likely to comply [117]. In addition, students-as-partners in AI governance may also have the indirect benefit of improving relationships between students and instructors [23, 86]. Given existing breakdowns of bi-directional trust among students and instructors [101], participatory approaches to AI governance may help to repair breakdowns by calling students *in* rather than calling students *out*.

We can also understand students’ participation in the context the formation of *publics*, or, in other words, how zine-making and zine-sharing catalyzed participation which persisted beyond the bounds of the workshop [34]. In our study, zines created conditions for such a public to take shape, comprising not only students but those concerned about the lack of student voice in AI governance: recall the several instructors who, after flipping through the zine, began to explore ways to solicit student voice in AI policy. Zine-making in particular, was an important artifact for this public to take shape because institutional feedback mechanisms—surveys, comment forms, office hours—often function to absorb concerns without addressing them [11]; evidenced by the plethora of large-scale surveys soliciting student perspectives on AI [125], and very little evidence that these data are translated into governance changes [82]. Instead, our zine-based approach sought to create conditions where students could articulate discontent and concerns outside these traditional channels.

5.2.3 Recommendations for Fostering Participatory AI Governance in a Classroom near You. We understand that zine-making is not always feasible or desirable. Other approaches to foster student participation include campus-, department-, and course-level interventions. For larger-scale participation projects, interventions such as listening sessions or town halls may be ideal, especially when interdisciplinary discourse on AI policy may be desired. For course-level interventions, instructors could begin class with a co-articulation of policies and agreements, and then solicit students’

formative feedback midway through the term on how AI policies are and are not working (and why). Policies can be embedded into other artifacts, besides zines, better suited to different disciplines: one-pagers for professional programs, podcast episodes for communications students, GitHub repositories for computer science students, and so on.

Across these different approaches, however, equitable participation cannot be presumed—it must be actively cultivated through attention to power dynamics [10, 37, 52]. Throughout this project, creating student-centered spaces was not simply a methodological preference but a prerequisite for surfacing concerns that other approaches miss [101]. This means attending to who facilitates (student-led or neutral bodies like university libraries may elicit more candid responses than faculty-led effort), how data is collected (anonymous approaches reduce fear of traceability), and whether the process is ongoing. On this last point: policies are not one and done. The AI landscape changes quickly as new tools and capabilities emerge, and policies wedged to particular platforms or interaction patterns (e.g., chat-based dialogue systems) may become deprecated rapidly. It is therefore essential to treat policies as living documents, creating sustained channels of communication with students rather than one-off consultations.

6 Limitations

This study draws on a single graduate design cohort at a minority-serving public university, limiting generalizability. Facilitators seeded policy topics as sensitizing entry points to AI policy design [108], but this may have induced confirmation bias. In addition, this study centers student perspectives—by design—and does not incorporate educator or administrator viewpoints. The resulting policies therefore are not always feasible or reasonable to implement, especially when incongruent with learning objectives. Finally, our workshops did not address consequences for policy violations, but rather explored how student participation may lessen the likelihood for violations.

7 Conclusion

In this paper, we sought student perspectives on AI policies in a design classroom, towards participatory governance of AI in higher education. To do so, we positioned students as lead users, or early adopters of generative AI technologies, and looked to their lived experience with these tools as a source of legitimate information to inform policies. Through a three-part workshop series, students authored ten policy recommendations, embedded in a zine that has since circulated across campus, sparking discourse around student participation beyond the bounds of the workshops. These recommendations surfaced concerns often absent from top-down AI governance—such as the desire for assignment-specific guidance, frustration with perceived double standards in faculty AI use, and support for English learners—though they also raised implementation challenges, particularly when student preferences diverged from learning objectives. Finally, we argue that while the policies themselves are specific to a design classroom, the methods for fostering student participation are transferable to other domains. We offer strategies for creating channels for student input that

instructors across disciplines can adapt to their own contexts. As generative AI continues to reshape higher education, we hope this work offers one model for calling students *in* rather than calling students *out*.

Acknowledgments

This work was supported by UMBC's Hrabowski Pedagogy Innovation Award. First, we thank our participants for their deep and meaningful contributions to this work. We also thank Julie Sayo and Aasmita Bhattacharya for zine guidance and designerly excellence, and Lara Martin for feedback and support. Finally, we thank our reviewers for their encouragement and constructive feedback. Generative AI tools like ChatGPT and Claude were used to assist with related work searches and summarization of text.

References

- [1] BoodleBox [n. d.]. *BoodleBox*. BoodleBox. <https://boodlebox.ai/>
- [2] OpenAI [n. d.]. *ChatGPT*. OpenAI. <https://chatgpt.com/>
- [3] Anthropic [n. d.]. *Claude*. Anthropic. <https://claude.ai/login>
- [4] OpenAI [n. d.]. *DALL-E 3*. OpenAI. <https://openai.com/index/dall-e-3/>
- [5] Figma Community [n. d.]. *Iconify - Figma Plugin*. Figma Community. <https://www.figma.com/community/plugin/735098390272716381/iconify>
- [6] Uizard Technologies [n. d.]. *Uizard: AI-Powered Design Tool*. Uizard Technologies. <https://uizard.io/>
- [7] [n. d.]. *UX Pilot AI: Design Apps with AI*. <https://uxpilot.ai/>
- [8] Vercel, Inc. [n. d.]. *v0 by Vercel*. Vercel, Inc. <https://v0app/>
- [9] Nathalie Aelterman, Maarten Vansteenkiste, and Leen Haerens. 2019. Correlates of students' internalization and defiance of classroom rules: A self-determination theory perspective. *British journal of educational psychology* 89, 1 (2019), 22–40.
- [10] Nimra Ahmed, Luise Arn, Natalia Obukhova, Nicolai Brodersen Hansen, Nadia Campo Woytuk, Anupriya Tuli, Karin Hansson, Angelika Strohmayer, and Elaine M Huang. 2025. Ethics, Power, and Tensions: Rethinking Participation-Based Design for Sensitive Contexts. In *Adjunct Proceedings of the Sixth Decennial Aarhus Conference: Computing X Crisis*. 1–4.
- [11] Sara Ahmed. 2021. *Complaint!* Duke University Press.
- [12] Rami Alshareefen and Naji Al Sayari. 2025. Examining academic integrity policy and practice in the era of AI: a case study of faculty perspectives. *Frontiers in Education* 10 (2025). doi:10.3389/feduc.2025.1621743
- [13] Yunjo An, Ji Hyun Yu, and Shadarra James. 2025. Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education* 22, 1 (2025), 10.
- [14] Mariam Asad and Christopher A. Le Dantec. 2017. Tap the "Make This Public" Button: A Design-Based Inquiry into Issue Advocacy and Digital Civics. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 6304–6316. doi:10.1145/3025453.3026034
- [15] Chris Atton. 1999. A reassessment of the alternative press. *Media, Culture & Society* 21, 1 (1999), 51–76. doi:10.1177/016344399021001003
- [16] Alex Barrett and Austin Pack. 2023. Not quite eye to AI: Student and teacher perspectives on the use of generative artificial intelligence in the writing process. *International Journal of Educational Technology in Higher Education* 20, 1 (2023), 59.
- [17] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Jason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–8.
- [18] David Birks and John Clare. 2023. Linking artificial intelligence facilitated academic misconduct to existing prevention frameworks. *International Journal for Educational Integrity* 19, 20 (2023), 1–15. doi:10.1007/s40979-023-00142-3
- [19] Elizabeth L Bjork, Robert A Bjork, et al. 2011. Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. *Psychology and the real world: Essays illustrating fundamental contributions to society* 2, 59–68 (2011), 56–64.
- [20] Claus Bossen, Christian Dindler, and Ole Sejer Iversen. 2022. Challenges of Scaling Participatory Design: A Systematic Literature Review. In *Proceedings of the 34th Australian Conference on Human-Computer Interaction (OzCHI '22)*. ACM. doi:10.1145/3572921.3572924
- [21] Isabelle Bousquette. 2024. The Hottest AI Job of 2023 Is Already Obsolete. *The Wall Street Journal* (April 25 2024). <https://www.wsj.com/articles/the-hottest-ai-job-of-2023-is-already-obsolete-1961b054> Accessed: February 4, 2026.
- [22] Catherine Bovill. 2020. Co-creation in learning and teaching: The case for a whole-class approach in higher education. *Higher Education* 79 (2020), 1023–1037. doi:10.1007/s10734-019-00453-w
- [23] Catherine Bovill, Alison Cook-Sather, and Peter Felten. 2011. Students as co-creators of teaching approaches, course design, and curricula: Implications for academic developers. *International Journal for Academic Development* 16, 2 (2011), 133–145. doi:10.1080/1360144X.2011.568690
- [24] José Antonio Bowen and C. Edward Watson. 2024. *Teaching with AI*. Johns Hopkins University Press. doi:10.56021/9781421449227
- [25] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [26] Virginia Braun and Victoria Clarke. 2022. Thematic analysis: a practical guide. Thematic analysis.
- [27] Brown University Library. 2026. Citation and Attribution - Generative Artificial Intelligence. <https://libguides.brown.edu/c.php?g=1338928&p=9868287> Last updated January 29, 2026; Accessed: February 2, 2026.
- [28] Alice Cai, Ian Arawjo, and Elena L Glassman. 2024. Antagonistic ai. *arXiv preprint arXiv:2402.07350* (2024).
- [29] Cecilia Ka Yuk Chan. 2023. A Comprehensive AI Policy Education Framework for University Teaching and Learning. *International Journal of Educational Technology in Higher Education* 20, 1 (July 2023), 38. doi:10.1186/s41239-023-00408-3
- [30] Kathy Charmaz. 2006. *Constructing grounded theory: A practical guide through qualitative analysis*. sage.
- [31] Liuqing Chen, Yaxuan Song, Jia Guo, Lingyun Sun, Peter Childs, and Yuan Yin. 2025. How Generative AI Supports Human in Conceptual Design. *Design Science* 11 (Jan. 2025), e9. doi:10.1017/dsj.2025.2
- [32] Alison Cook-Sather, Catherine Bovill, and Peter Felten. 2014. *Engaging students as partners in learning and teaching: A guide for faculty*. John Wiley & Sons.
- [33] Cornell University English Language Support Office. 2025. AI and Academic Integrity. <https://elsos.as.cornell.edu/ai-and-academic-integrity> Developed September 2025; Accessed February 2, 2026.
- [34] Christopher A Le Dantec and Carl DiSalvo. 2013. Infrastructuring and the Formation of Publics in Participatory Design. *Social Studies of Science* 43, 2 (April 2013), 241–264. doi:10.1177/0306312712471581
- [35] Karl de Fine Licht. 2024. Generative artificial intelligence in higher education: Why the banning approach to student use is sometimes morally justified. *Philosophy & Technology* 37, 3 (2024), 113.
- [36] Edward L Deci, Haleh Eghrari, Brian C Patrick, and Dean R Leone. 1994. Facilitating internalization: The self-determination theory perspective. *Journal of personality* 62, 1 (1994), 119–142.
- [37] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23)*. Association for Computing Machinery, New York, NY, USA, 1–23. doi:10.1145/3617694.3623261
- [38] Stephen Duncombe. 1997. *Notes From Underground: Zines and the Politics of Alternative Culture*. Microcosm Publishing, Cleveland, OH.
- [39] Nagalakshmi Eragamreddy. 2025. From Plagiarism to Paraphrasing: Graduate Students' Approaches for Referencing Materials in Academic Writing. *Teaching English Language* 19, 2 (2025), 89–137. doi:10.22132/TEL.2024.232821
- [40] Gökberk Erol, Anıl Ergen, Büşra Gülsen Erol, Şebnem Kaya Ergen, Tevfik Serhan Bora, Ali Deniz Çölgeçen, Büşra Araz, Cansel Şahin, Günsu Bostancı, İlday Küçük, et al. 2025. Can we trust academic AI detective? Accuracy and limitations of AI-output detectors. *Acta neurochirurgica* 167, 1 (2025), 1–12.
- [41] Figma, Inc. 2025. Figma [Online collaborative design platform]. <https://www.figma.com/> Software.
- [42] Financial Times. 2024. Science students lead use of AI tools in assessments. <https://www.ft.com/content/d591fb1a-9fc6-4345-b5fc-781e091ae3f8>. Accessed: 2025-09-11.
- [43] N. J. Francis, S. Jones, and D. P. Smith. 2025. Generative AI in Higher Education: Balancing Innovation and Integrity. *British Journal of Biomedical Science* 81 (2025), 14048. doi:10.3389/bjbs.2024.14048
- [44] Aakash Gautam. 2024. Reconfiguring participatory design to resist ai realism. In *Proceedings of the Participatory Design Conference 2024: Exploratory Papers and Workshops-Volume 2*. 31–36.
- [45] Michael Gerlich. 2025. AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies* 15, 1 (2025), 6.
- [46] Aashish Ghimire and John Edwards. 2024. From Guidelines to Governance: A Study of AI Policies in Education. *arXiv:2403.15601 [cs.CY]* <https://arxiv.org/abs/2403.15601>
- [47] Levent Giray. 2024. The Problem with False Positives: AI Detection Unfairly Accuses Scholars of AI Plagiarism. *The Serials Librarian* 85, 5–6 (2024), 181–189. doi:10.1080/0361526X.2024.2433256
- [48] Chahna Gonsalves. 2025. Addressing student non-compliance in AI use declarations: implications for academic integrity and assessment in higher education. *Assessment & Evaluation in Higher Education* 50, 4 (2025), 592–606.

[49] Tim Gorichanaz. 2023. Accused: How students respond to allegations of using ChatGPT on assessments. *Learning: Research and Practice* 9, 2 (2023), 183–196.

[50] Grammarly, Inc. 2025. Grammarly [AI writing assistance software]. <https://www.grammarly.com/> Software.

[51] Bingyi Han, Sadia Nawaz, George Buchanan, and Dana McKay. 2025. Students' Perceptions: Exploring the Interplay of Ethical and Pedagogical Impacts for Adopting AI in Higher Education. *International Journal of Artificial Intelligence in Education* (Jan. 2025). doi:10.1007/s40593-024-00456-4

[52] Christina Harrington, Sheena Erete, and Anna Marie Piper. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–25.

[53] Kiersten Hay. 2022. Zineography: Exploring the participatory design process of collaborative zine making. In *Proceedings of the Participatory Design Conference 2022-Volume 2*. 313–316.

[54] Kiersten Hay, Abigail C Durrant, Shema Tariq, Lynne Coventry, and Helen Anderson. 2024. Zineography: A Community-Based Research-through-Design Method of Zine Making for Unequal Contexts. In *Nordic Conference on Human-Computer Interaction*. ACM, Uppsala Sweden, 1–17. doi:10.1145/3679318.3685390

[55] Mick Healey. 2014. Students as partners in learning and teaching in higher education. In *Workshop Presented at University College Cork*, Vol. 12. 15.

[56] Gerald F Hess. 2007. Collaborative course design: Not my course, not their course, but our course. *Washburn LJ* 47 (2007), 367.

[57] Catherine M Hicks, Vineet Pandey, C Ailie Fraser, and Scott Klemmer. 2016. Framing feedback: Choosing review environment features that support high quality peer assessment. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 458–469.

[58] Kashmir Hill. 2025. The Professors Are Using ChatGPT, and Some Students Aren't Happy About It. *The New York Times* (May 14 2025). <https://www.nytimes.com/2025/05/14/technology/chatgpt-college-professors.html> Technology section.

[59] Marieta Hristova. 2025. Generative Artificial Intelligence and Academic Practices: A Comparative Analysis of Approaches in Europe, the United States and China. *Postmodernism Problems* 15, 2 (2025), 245–269. doi:10.46324/PMP2502245

[60] Hsiu-Ling Hsiao and Hsien-Hui Tang. 2024. A Study on the Application of Generative AI Tools in Assisting the User Experience Design Process. In *Artificial Intelligence in HCI*. Helmut Degen and Stavroula Ntoa (Eds.). Springer Nature Switzerland, Cham, 175–189. doi:10.1007/978-3-031-60611-3_13

[61] Yueqiao Jin, Lixiang Yan, Vanessa Echeverria, Dragan Gašević, and Roberto Martínez-Maldonado. 2025. Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence* 8 (2025), 100348.

[62] Heather Johnston, Rebecca F. Wells, Elizabeth M. Shanks, Timothy Boey, and Bryony N. Parsons. 2024. Student Perspectives on the Use of Generative Artificial Intelligence Technologies in Higher Education. *International Journal for Educational Integrity* 20, 1 (Dec. 2024), 1–21. doi:10.1007/s40979-024-00149-4

[63] René F. Kizilcec, Elaine Huber, Elena C. Papanastasiou, Andrew Cram, Christos A. Makridis, Adele Smolansky, Sandris Žeivots, and Corina Radulescu. 2024. Perceived Impact of Generative AI on Assessments: Comparing Educator and Student Perspectives in Australia, Cyprus, and the United States. *Computers and Education: Artificial Intelligence* 7 (Dec. 2024), 100269. doi:10.1016/j.caei.2024.100269

[64] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Patti Maes. 2025. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. *arXiv preprint arXiv:2506.08872* 4 (2025).

[65] Yasmine Kotturi, Angel Anderson, Glenn Ford, Michael Skirpan, and Jeffrey P Bigham. 2024. Deconstructing the Veneer of Simplicity: Co-Designing Introductory Generative AI Workshops with Local Entrepreneurs. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3613904.3642191

[66] Chinmay E Kulkarni, Michael S Bernstein, and Scott R Klemmer. 2015. Peer Studio: rapid peer feedback emphasizes revision and improves performance. In *Proceedings of the second (2015) ACM conference on learning@ scale*. 75–84.

[67] Rahul Kumar. 2023. Faculty members' use of artificial intelligence to grade student papers: A case of implications. *International Journal for Educational Integrity* 19, 1 (2023), 9.

[68] Tzu-Sheng Kuo, Quan Ze Chen, Amy X. Zhang, Jane Hsieh, Haiyi Zhu, and Kenneth Holstein. 2025. PolicyCraft: Supporting Collaborative and Participatory Policy Design through Case-Grounded Deliberation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–24. arXiv:2409.15644 [cs] doi:10.1145/3706598.3713865

[69] Lin Kyi, Amruta Mahuli, M Six Silberman, Reuben Binns, Jun Zhao, and Asia J Biega. 2025. Governance of Generative AI in Creative Work: Consent, Credit, Compensation, and Beyond. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.

[70] LanguaTalk. 2021. LanguaTalk: Learn a Language with the World's Best Tutors and AI. Online Platform. <https://languatalk.com/> Accessed: February 4, 2026.

[71] Bruno Latour. 2008. A cautious Prometheus? A few steps toward a philosophy of design. In *"Networks of Design"*, Annual International Conference of the Design History Society. Universal Publishers, 2–10.

[72] Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–22.

[73] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User experience design professionals' perceptions of generative artificial intelligence. In *Proceedings of the 2024 CHI conference on human factors in computing systems*. 1–18.

[74] Adela C Licona. 2012. *Zines in third space: Radical cooperation and borderlands rhetoric*. Suny Press.

[75] Fang Liu, Junyan Lv, Shenglan Cui, Zhilong Luan, Kui Wu, and Tongqing Zhou. 2024. Smart "Error"! Exploring Imperfect AI to Support Creative Ideation. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1 (April 2024), 121:1–121:28. doi:10.1145/3637398

[76] Jiahui Luo. 2024. A critical review of GenAI policies in higher education assessment: A call to reconsider the "originality" of students' work. *Assessment & Evaluation in Higher Education* 49, 5 (2024), 651–664.

[77] Jiahui Luo. 2025. How does GenAI affect trust in teacher-student relationships? Insights from students' assessment experiences. *Teaching in Higher Education* 30, 4 (2025), 991–1006.

[78] Frank Lyman. 1981. The responsive classroom discussion: The inclusion of all students. *Mainstreaming digest*. College Park, MD: University of Maryland (1981).

[79] Stephen Marshall, Lia Blaj-Ward, Namhee Dreamson, Joan Nyanjom, and Maria T. Bertuol. 2024. The reshaping of higher education: technological impacts, pedagogical change, and future projections. *Higher Education Research & Development* 43, 3 (2024), 521–541. doi:10.1080/07294360.2024.2329393

[80] Kelly Matthews. 2016. Students as partners as the future of student engagement. *Student Engagement in Higher Education Journal* 1, 1 (2016).

[81] Kelly E Matthews. 2018. Engaging students as participants and partners: An argument for partnership with students in higher education research on student success. *International Journal of Chinese Education* 7, 1 (2018), 42–64.

[82] Nora McDonald, Aditya Johri, Areej Ali, and Aayushi Hingle Collier. 2025. Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines. *Computers in Human Behavior: Artificial Humans* 3 (March 2025), 100121. doi:10.1016/j.chbh.2025.100121

[83] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.

[84] Megan McNamara. 2025. Sociology 136: Social Psychology [Course Syllabus]. Summer Session 1. https://summer.ucsc.edu/files/2025/10/2025_socy_136_mcnamara.pdf

[85] Andrew McNutt. 2021. On the potential of zines as a medium for visualization. In *2021 IEEE Visualization Conference (VIS)*. IEEE, 176–180.

[86] Lucy Mercer-Mapstone, Sofia Louise Dvorakova, Kelly E. Matthews, Sophia Abbot, Breagh Cheng, Peter Felten, Kali Knorr, Elizabeth Marquis, Rafaela Shamber, and Kelly Swaim. 2017. A systematic literature review of students as partners in higher education. *International Journal for Students as Partners* 1, 1 (2017). doi:10.15173/ijssap.v1i1.3119

[87] Miro. 2025. Miro [Online collaborative whiteboard platform]. <https://miro.com/> Software.

[88] Ethan Mollick and Lilach Mollick. 2023. AI as a Team Coach. *Harvard Business Publishing* (2023). <https://hbsp.harvard.edu/inspiring-minds/ai-as-team-coach> Published on September 25, 2023.

[89] Benjamin Luke Moorhouse, Marie Alina Yeo, and Yuwei Wan. 2023. Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open* 5 (2023), 100151.

[90] Isabel Moreno-Lopez. 2005. Sharing power with students: The critical language classroom. *Radical Pedagogy* 7, 2 (2005), 23–49.

[91] Michael J Muller and A Druin. 2007. Participatory design: The third space in HCI (revised). *Handbook of HCI 2nd Edition*. Mahwah NJ USA: Erlbaum (2007).

[92] Ileana Najarro. 2023. *What Teachers Should Know Before Using AI With English Learners*. <https://www.edweek.org/teaching-learning/what-teachers-should-know-before-using-ai-with-english-learners/2023/10>

[93] Ohio State University. 2025. *Ohio State launches bold AI Fluency initiative to redefine learning and innovation*. <https://news.osu.edu/ohio-state-launches-bold-ai-fluency-initiative-to-redefine-learning-and-innovation> Accessed: 2026-01-30.

[94] Jonas Oppenlaender, Rhema Linder, and Johanna Silvennoinen. 2025. Prompting AI art: An investigation into the creative skill of prompt engineering. *International journal of human-computer interaction* 41, 16 (2025), 10207–10229.

[95] Otter.ai, Inc. 2025. Otter.ai [AI notetaking and transcription software]. <https://otter.ai/> Software.

[96] Weina Pang and Zhe Wei. 2025. Shaping the future of higher education: A technology usage study on generative AI innovations. *Information* 16, 2 (2025), 95.

[97] PatorJK, [n. d.]. *ASCII Text Art Generator*. <https://patorjk.com/software/taag/#p=display&f=Graffiti&t=Type+Something+&x=none&v=4&h=4&w=80&w=false>

[98] Tiffany Petricini, Sarah Zipf, and Chuhao Wu. 2025. RESEARCH-AI: Communicating academic honesty: teacher messages and student perceptions about generative AI. *Frontiers in Communication* 10 (jun 2025). doi:10.3389/fcomm.2025.1544430

[99] Jessie Pink. 2016. Tracing the beginnings of the punk fanzine. *Dazed* (May 6 2016). <https://www.dazeddigital.com/artsandculture/article/30999/1/mark-perry-tracing-the-beginnings-of-the-punk-fanzine> Accessed: YYYY-MM-DD.

[100] Tejaswi Polimetta and Katy Ilonka Gero. 2025. A Paradigm for Creative Ownership. *arXiv preprint arXiv:2505.15971* (2025).

[101] Isabella Pu, Prerna Ravi, Linh Dieu Dinh, Chelsea Joe, Caitlin Ogoe, Zixuan Li, Cynthia Breazeal, and Anastasia K. Ostrowski. 2025. "How Can We Learn and Use AI at the Same Time?": Participatory Design of GenAI with High School Students. In *Proceedings of the 24th Interaction Design and Children*. ACM, Reykjavik Iceland, 204–220. doi:10.1145/3713043.3727057

[102] Queer Music Heritage. n.d. Lisa Ben & Vice Versa. <https://queermusicheritage.com/viceversa.html> Accessed: 2025-09-11.

[103] Afshaneh Razi, Layla Bouzoubaa, Aria Pessianzadeh, John S. Seberger, and Rezvaneh (Shadi) Rezapour. 2025. Not a Swiss Army Knife: Academics' Perceptions of Trade-Offs around Generative AI Use. *Proceedings of the Association for Information Science and Technology* 62, 1 (Oct. 2025), 547–560. doi:10.1002/pra2.1277

[104] Justin Reich. 2025. Stop Pretending You Know How to Teach AI. *The Chronicle of Higher Education* (5 November 2025). <https://www.chronicle.com/article/stop-pretending-you-know-how-to-teach-ai> Opinion essay. A version appeared in the November 14, 2025 issue.

[105] Jose Rosa. 2025. Artificial Intelligence and the Future of Personalized Learning: The Case of Arizona State University. doi:10.13140/RG.2.2.25056.96003

[106] Rajdeep Sah, Cameron Hagemaster, Arpan Adhikari, Ari Lee, and Nan Sun. 2025. Generative AI in higher education: student and faculty perspectives on use, ethics, and impact. *Issues in Information Systems* 26, 2 (2025).

[107] Sue Samson. 2010. Information literacy learning outcomes and student success. *The Journal of Academic Librarianship* 36, 3 (2010), 202–210.

[108] Elizabeth B-N Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *Co-design* 4, 1 (2008), 5–18.

[109] Advait Sarkar. 2023. Enough with "human-AI collaboration". In *Extended abstracts of the 2023 CHI conference on human factors in computing systems*. 1–8.

[110] Arvind Satyanarayan and G. Michael Jones. 2024. Intelligence as Agency: Evaluating the Capacity of Generative AI to Empower or Constrain Human Action. *An MIT Exploration of Generative AI* (2024). doi:10.21428/e4baedd9.2d7598a2

[111] Sandeep Singh Sengar, Affan Bin Hasan, Sanjay Kumar, and Fiona Carroll. 2025. Generative Artificial Intelligence: A Systematic Review and Applications. *Multimedia Tools and Applications* 84, 21 (June 2025), 23661–23700. doi:10.1007/s11042-024-20016-1

[112] Catherine Shaw, Li Yuan, Daniel Brennan, Sterling Martin, Natasha Johnson, Kristen Fox, and Gates Bryant. 2023. *GenAI in Higher Education: Fall 2023 Update Time for Class Study*. Technical Report. Tyton Partners. <https://tytonpartners.com/app/uploads/2023/10/GenAI-IN-HIGHER-EDUCATION-FALL-2023-UPDATE-TIME-FOR-CLASS-STUDY.pdf>

[113] Ira Shor. 1996. *When students have power: Negotiating authority in a critical pedagogy*. University of Chicago Press.

[114] Michelle Smit, Reinhard F Wagner, and Taryn Jane Bond-Barnard. 2025. Ambiguous regulations for dealing with AI in higher education can lead to moral hazards among students. *Project Leadership and Society* (2025), 100187.

[115] Adele Smolansky, Andrew Cram, Corina Radulescu, Sandris Zeivots, Elaine Huber, and Rene F Kizilcec. 2023. Educator and student perspectives on the impact of generative AI on assessments in higher education. In *Proceedings of the tenth ACM conference on Learning@ Scale*. 378–382.

[116] Jane Southworth, Kati Migliaccio, Joe Glover, Ja'Net Glover, David Reed, Chris McCarty, Joel Brendemuhl, and Andrew Thomas. 2023. Developing a model for AI Across the Curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence* 4 (2023), 100127. doi:10.1016/j.caeari.2023.100127

[117] Rebecca R Steingut, Erika A Patall, and Scott S Trimble. 2017. The effect of rationale provision on motivation and performance outcomes: A meta-analysis. *Motivation Science* 3, 1 (2017), 19.

[118] StudentAIImpact. 2023. Develop Your Policy. <https://www.studentaiimpact.lmcgatech.edu/developing-your-policy/> Accessed February 1, 2026.

[119] Marielle Justine Sumilong. 2025. Instructional affect and learner motivation in generative AI-restrictive and permissive classrooms. In *Frontiers in Education*, Vol. 10. Frontiers Media SA, 1626802.

[120] Teaching and Learning Conestoga. 2025. Describe and Document AI Use. <https://tlconestoga.ca/describe-and-document-ai-use/> Faculty Learning Hub.

[121] The University of Queensland. 2023. EDUC3707: Teaching Humanities and Social Sciences Curriculum [Course Profile]. Semester 1, 2023. https://archive.course-profiles.uq.edu.au/student_section_loader/section_6/127696 Assessment section, AI tools policy.

[122] Maryam Tohidi, William Buxton, Ronald Baecker, and Abigail Sellen. 2006. Getting the Right Design and the Design Right. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. Association for Computing Machinery, New York, NY, USA, 1243–1252. doi:10.1145/1124772.1124960

[123] Jack Tsao. 2025. Trajectories of AI policy in higher education: Interpretations, discourses, and enactments of students and teachers. *Computers and Education: Artificial Intelligence* (2025), 100496.

[124] Turnitin. 2025. Turnitin. <https://www.turnitin.com> Accessed: 2025-12-03.

[125] Tyton Partners. 2025. Time for Class 2025: Empowering Educators, Engaging Students. https://4213961.fs1.hubspotusercontent-na1.net/hubfs/4213961/Publications/Time%20for%20Class/Tyton%20Partners_Time%20for%20Class%202025.pdf

[126] Harvard University. 2023. Guidelines for Using ChatGPT and Other Generative AI Tools at Harvard. Harvard University Website. <https://provost.harvard.edu/guidelines-using-chatgpt-and-other-generative-ai-tools-harvard> Office of the Provost.

[127] Liberty University. n.d. *Using Artificial Intelligence Responsibly*. Technical Report. Liberty University, Office of the Dean of Students. <https://www.liberty.edu/students/dean/wp-content/uploads/sites/91/Canvas-Announcement-for-Students.pdf> Canvas Announcement for Students.

[128] University Writing Program. 2026. Peer and AI Review + Reflection (PAIRR). <https://writing.ucdavis.edu/pairr> Last updated: January 28, 2026. Accessed: February 2, 2026.

[129] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring Participation: On How We Involve People in Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, 429–438. doi:10.1145/2470654.2470716

[130] Eric Von Hippel. 2006. *Democratizing innovation*. The MIT Press.

[131] Samangi Wadinambiarchachi, Ryan M. Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Vellosa. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3613904.3642919

[132] Qian Wan, Siying Hu, Yu Zhang, Piaohong Wang, Bo Wen, and Zhicong Lu. 2024. "It Felt Like Having a Second Mind": Investigating Human-AI Co-creativity in Prewriting with Large Language Models. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1 (April 2024), 84:1–84:26. doi:10.1145/3637361

[133] James Weichert, Dayoung Kim, Qin Zhu, Junghwan Kim, and Hoda Eldardiry. 2025. Assessing Computer Science Student Attitudes Towards AI Ethics and Policy. *AI and Ethics* (Aug. 2025). arXiv:2504.06296 [cs] doi:10.1007/s43681-025-00817-2

[134] Western Michigan University. 2024. AI and Critical Thinking in Education. Teaching and Learning Resources. <https://wmich.edu/x/teaching-learning/teaching-resources/ai-critical-thinking> Accessed: 2025-05-14.

[135] Hongbo Zhang, Pei Chen, Xuelong Xie, Chaoyi Lin, Lianyan Liu, Zhuoshu Li, Weitao You, and Lingyun Sun. 2024. ProtoDreamer: A Mixed-prototype Tool Combining Physical Model and Generative AI to Support Conceptual Design. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3654777.3676399

[136] Chengbo Zheng, Kangyu Yuan, Bingcan Guo, Reza Hadi Mogavi, Zhenhui Peng, Shuai Ma, and Xiaojuan Ma. 2024. Charting the Future of AI in Project-Based Learning: A Co-Design Exploration with Students. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3613904.3642807

[137] Jiayu Zheng, Lingxin Hao, Kelun Lu, Ashi Garg, Mike Reese, Melo-Jean Yap, I-Jeng Wang, Xingyun Wu, Wenrui Huang, Jenna Hoffman, et al. 2025. Do Students Rely on AI? Analysis of Student-ChatGPT Conversations from a Field Study. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, Vol. 8. 2796–2807.

[138] Kaitlyn Zhou, Kristina Gligorić, Myra Cheng, Michelle S Lam, Vyoma Raman, Boluwatife Aminu, Caeley Woo, Michael Brockman, Hannah Cha, and Dan Jurafsky. 2025. Attention to Non-Adopters. *arXiv preprint arXiv:2510.15951* (2025).

A Student-Driven AI Policy Recommendations Zine Pages



Figure A1: Instructions, Page 1 of 10

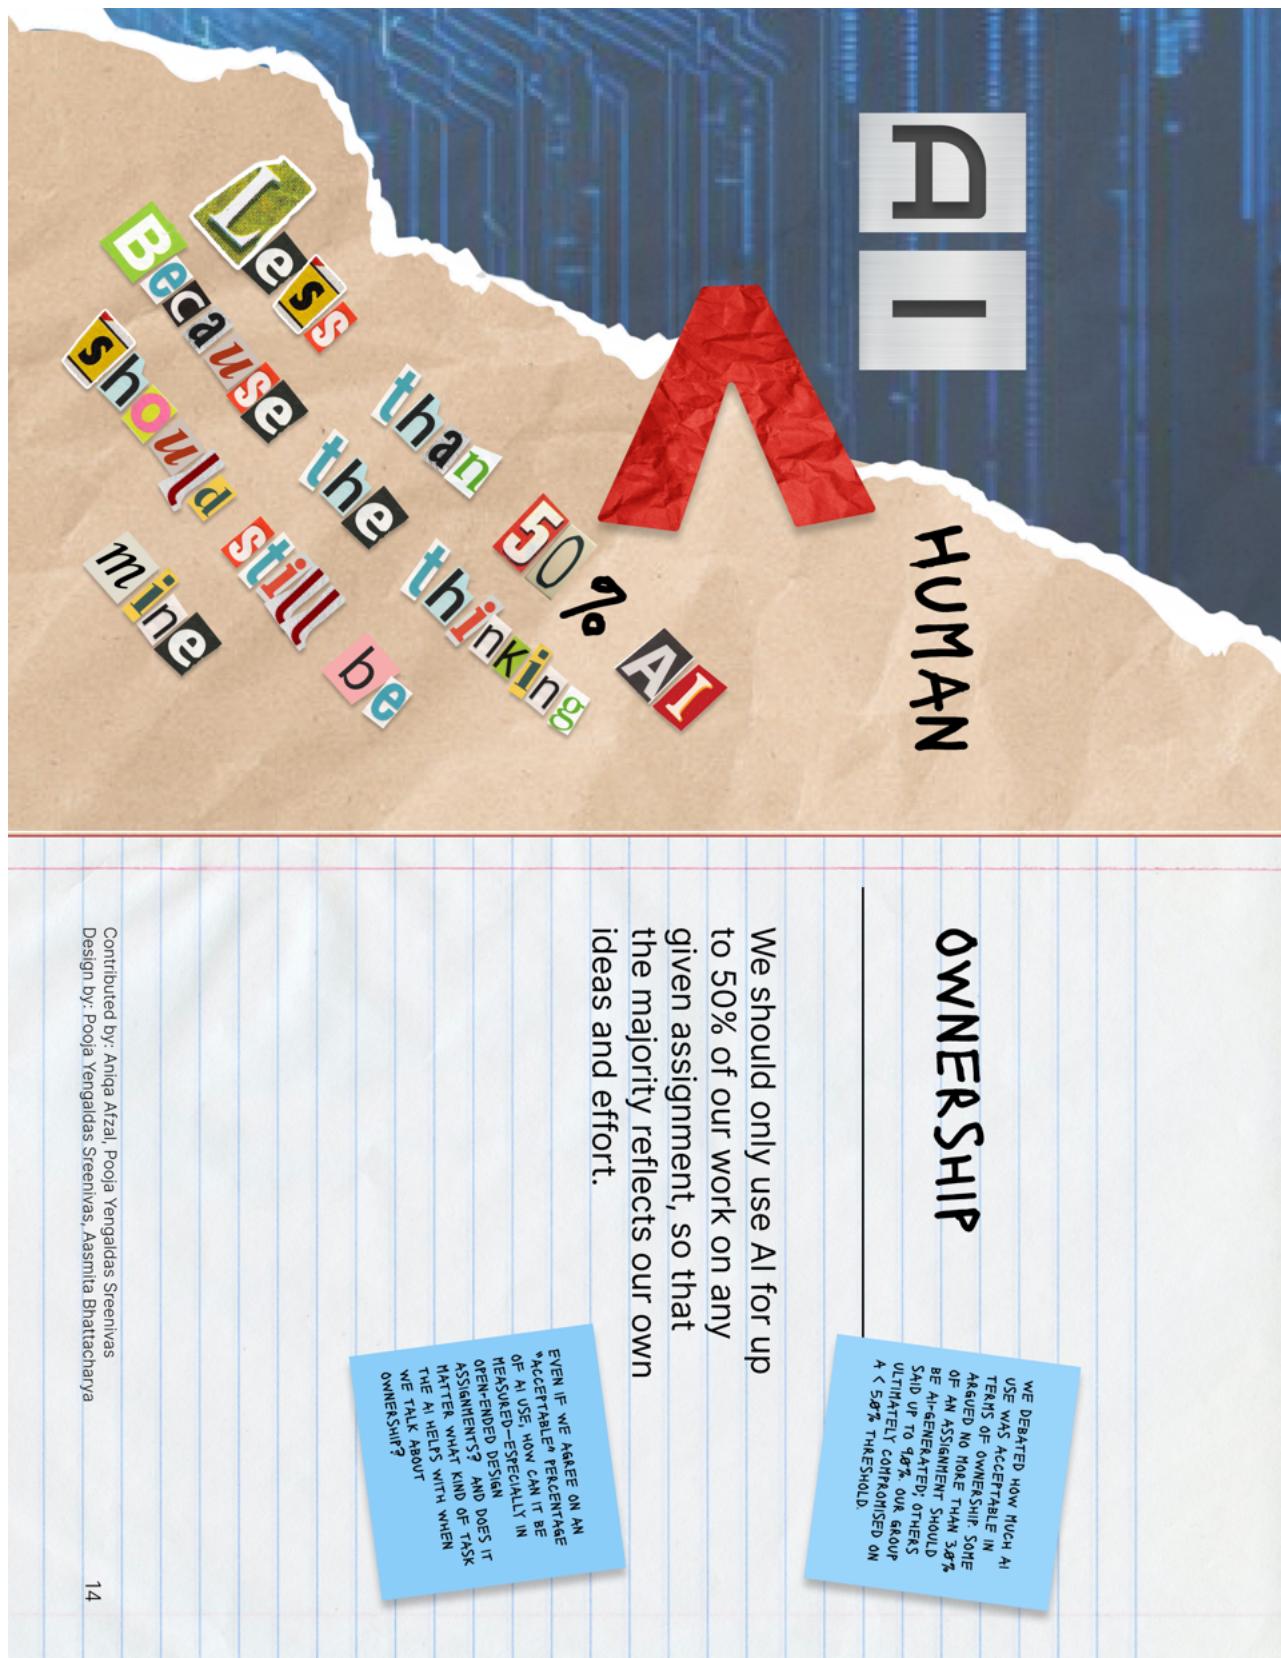


Figure A2: Ownership, Page 2 of 10

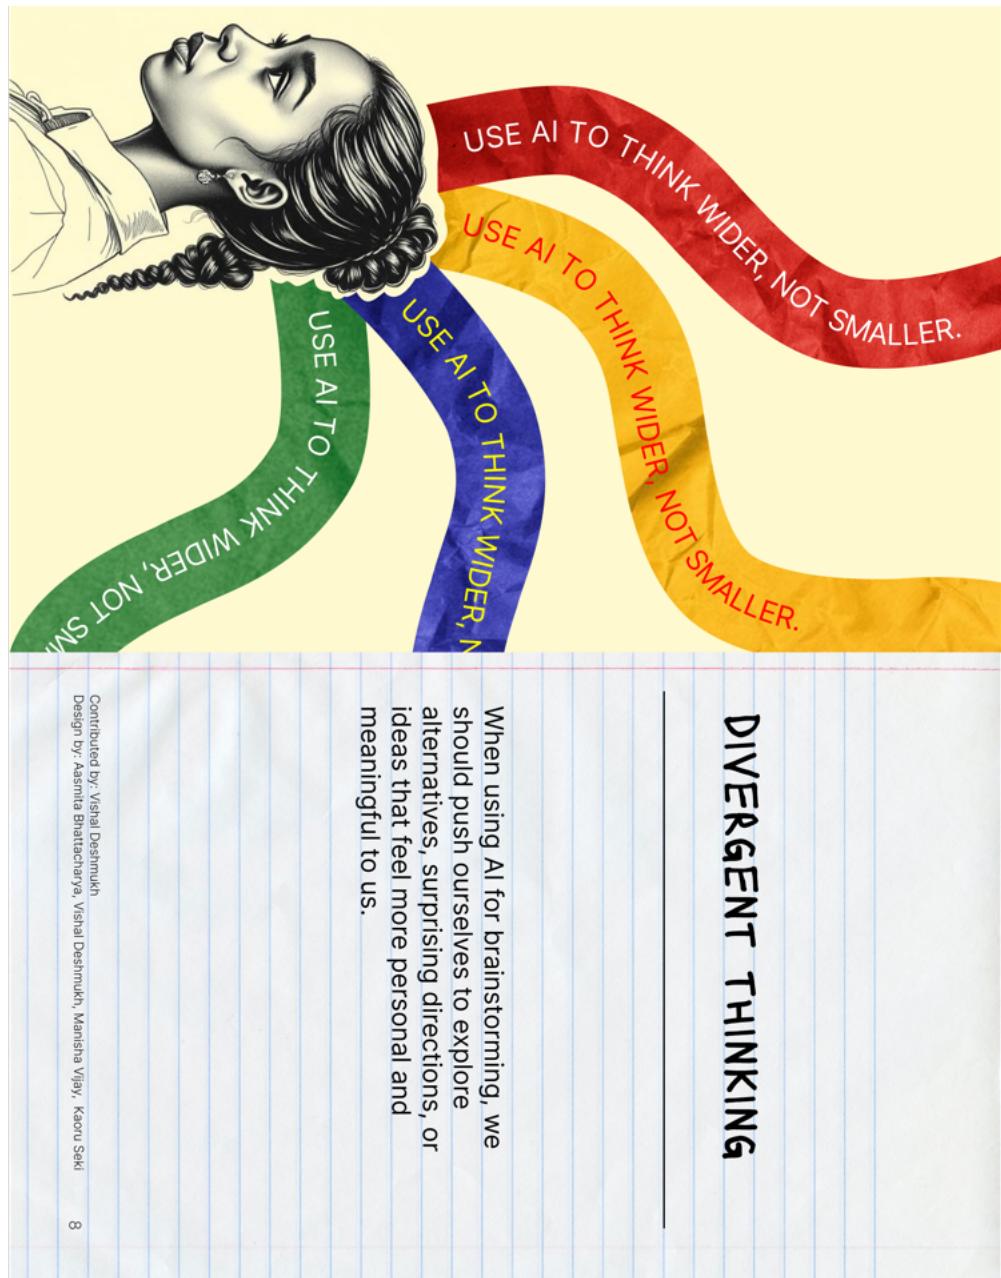
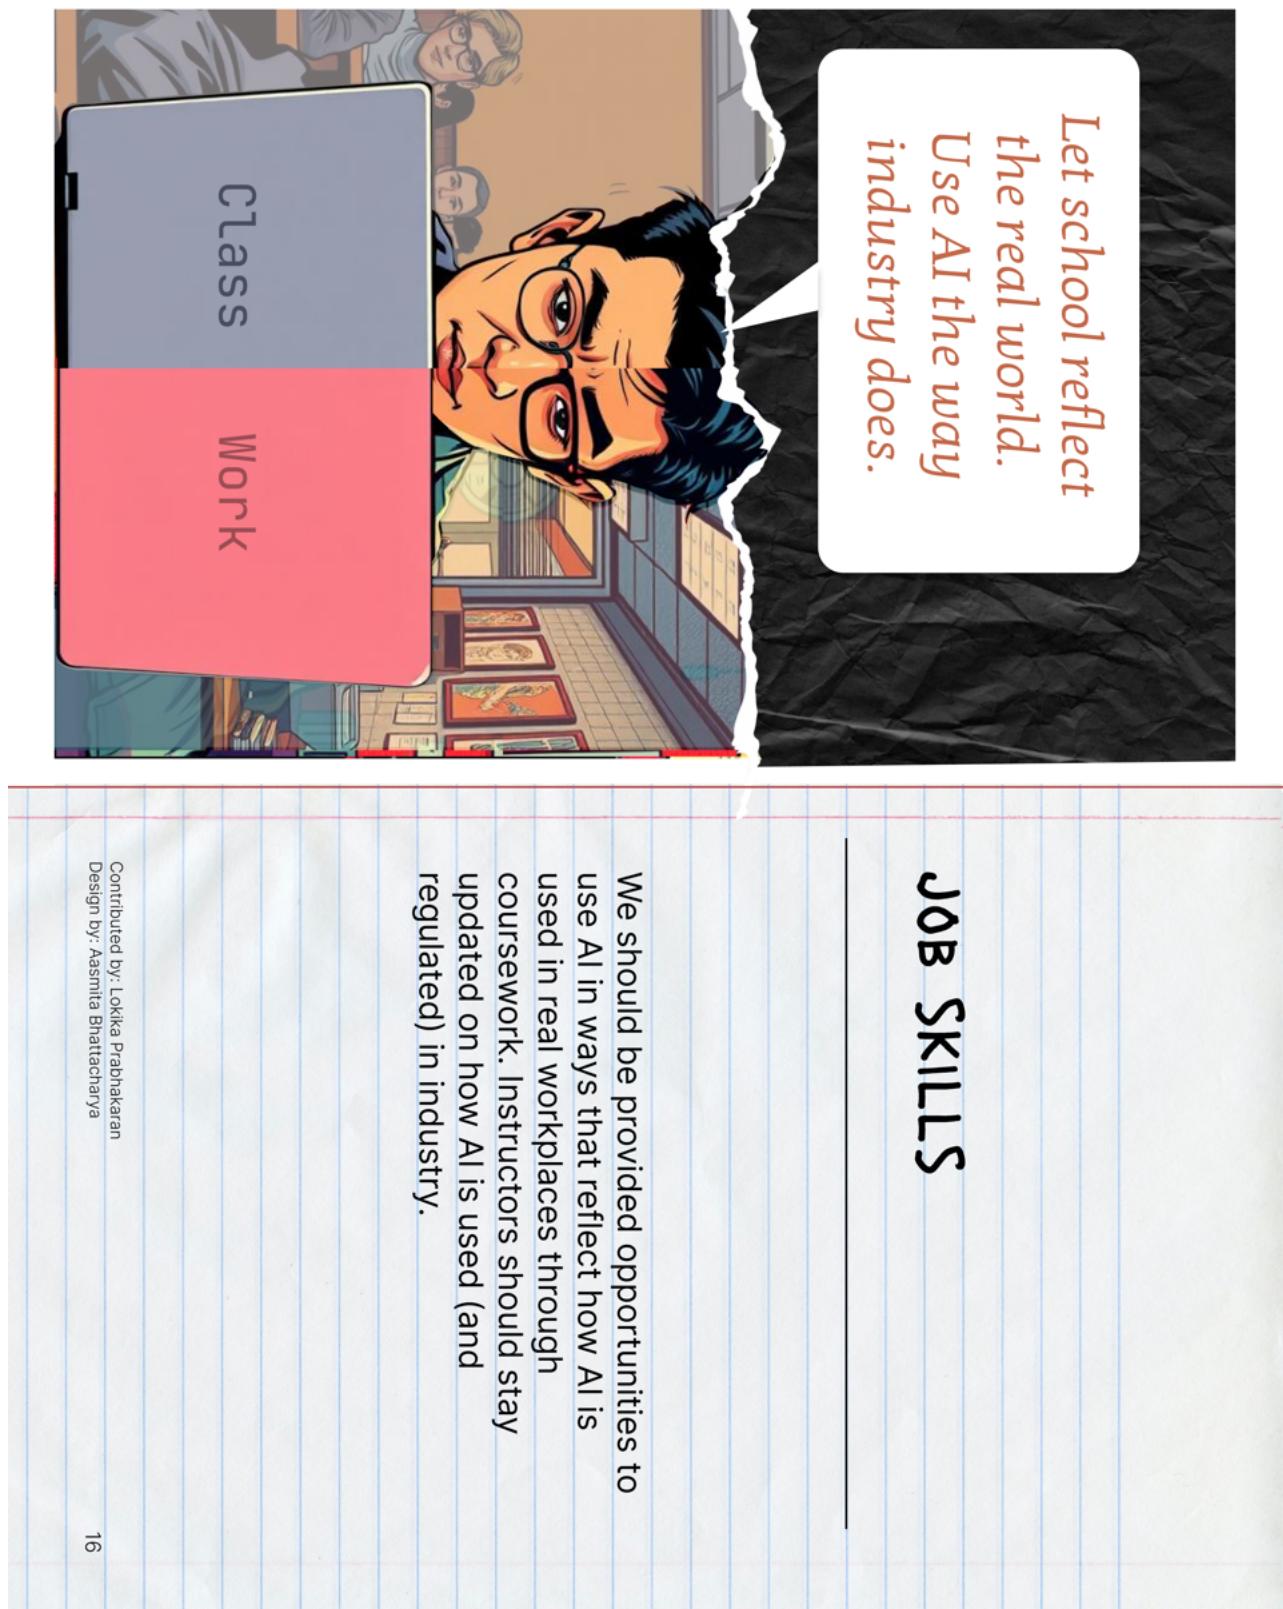


Figure A3: Divergent Thinking, Page 3 of 10



Let school reflect the real world. Use AI the way industry does.

JOB SKILLS

We should be provided opportunities to use AI in ways that reflect how AI is used in real workplaces through coursework. Instructors should stay updated on how AI is used (and regulated) in industry.

Contributed by: Lokika Prabhakaran
Design by: Aasmrita Bhattacharya

Figure A4: Job Skills, Page 4 of 10

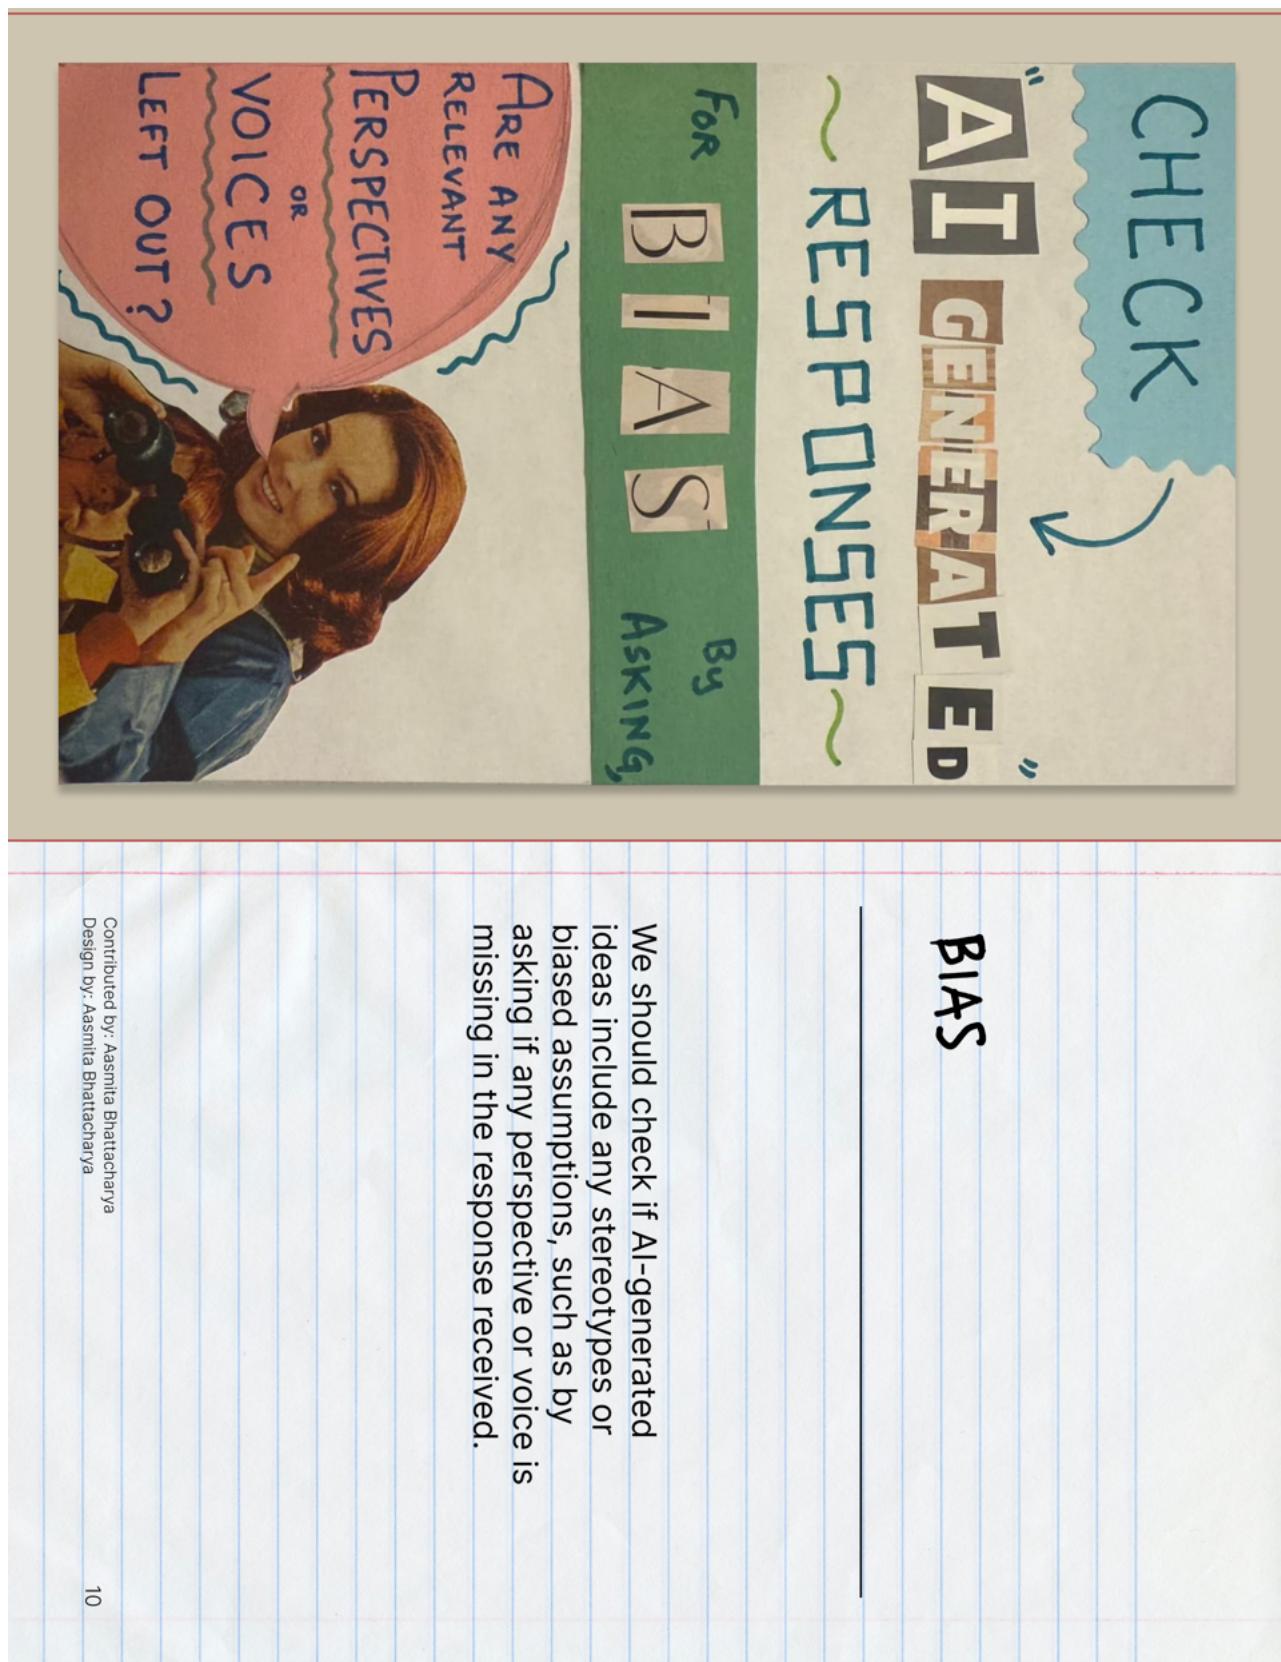


Figure A5: Bias, Page 5 of 10

GHOST WRITER.

AI ISN'T A

CITING AI USE

It is important for us to credit AI where credit is due. However, sharing chat logs, as many instructors currently require, is tedious and ineffective for both students and graders. Instead, we should share a 2-3 sentence summary for each submission describing how and why we used AI.

Contributed by: Aasmita Bhattacharya
Design by: Aasmita Bhattacharya

Figure A6: Citing AI Use, Page 6 of 10

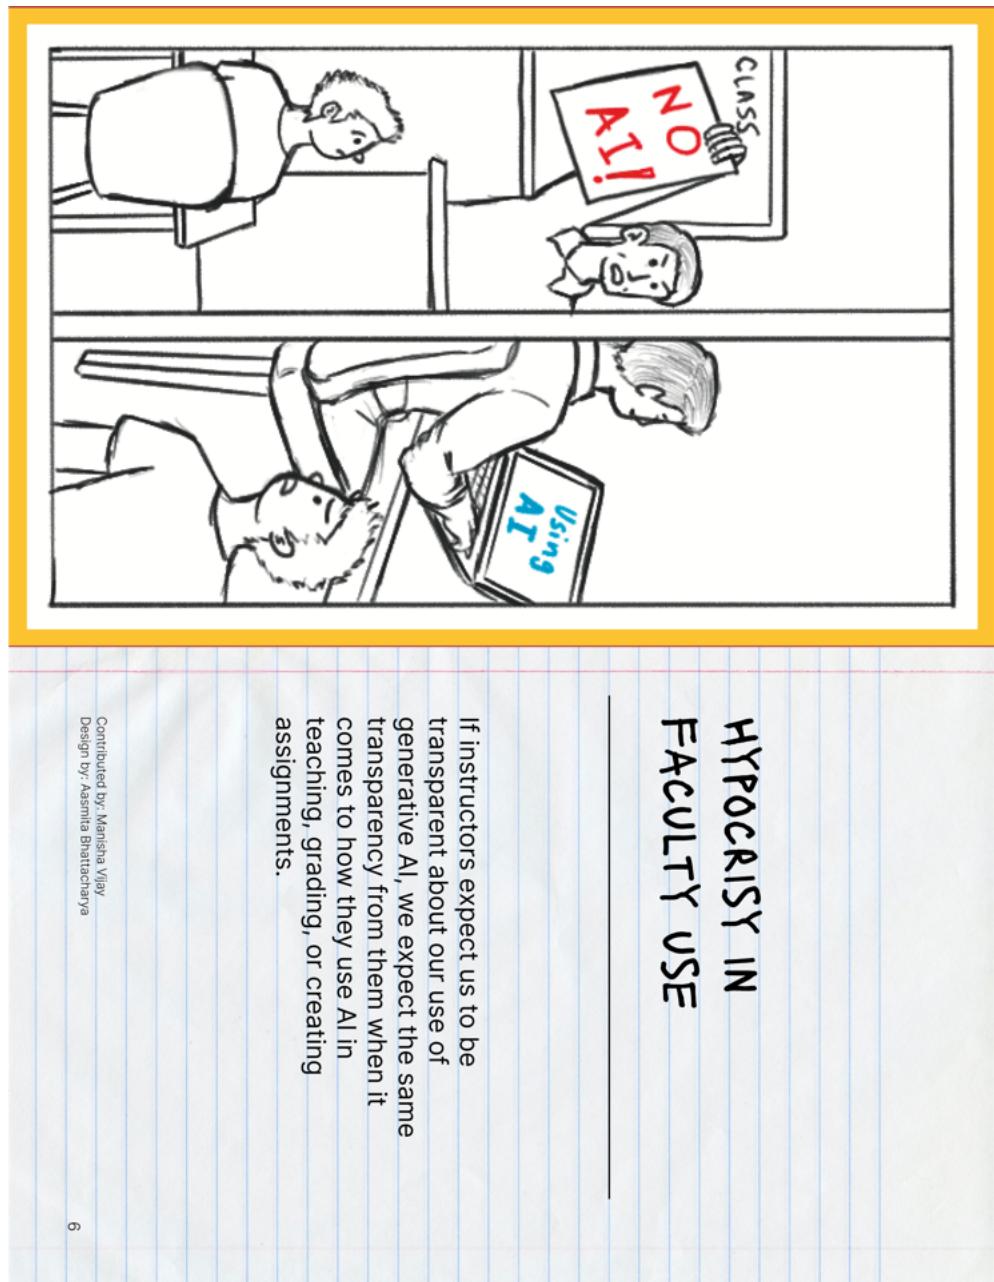


Figure A7: Hypocrisy in Faculty Use, Page 7 of 10

I feel the full
depth of meaning
of my thoughts,
but I can't always
convey these
feelings into
express them in
writing.

ENGLISH LEARNERS

For those of us who are English learners, we should be encouraged to use AI to support our English proficiency in writing by asking for refinement of our text. Also, we should ask for an explanation of the refinements made to further our language proficiency.

FEEDBACK

AI CAN TURN HARSH WORDS INTO HELPFUL ONES.

THIS MAKES NO SENSE

DO YOU EVEN KNOW WHAT YOU'RE TALKING ABOUT?

THIS IS JUST BAD

TERRIBLE ORGANIZATION

I think the structure can be refined more.

This section could use more contrast.

It feels a little underwhelming but I see the potential!

I DON'T LIKE IT

Clarify your main argument.

Figure A9: Feedback, Page 9 of 10

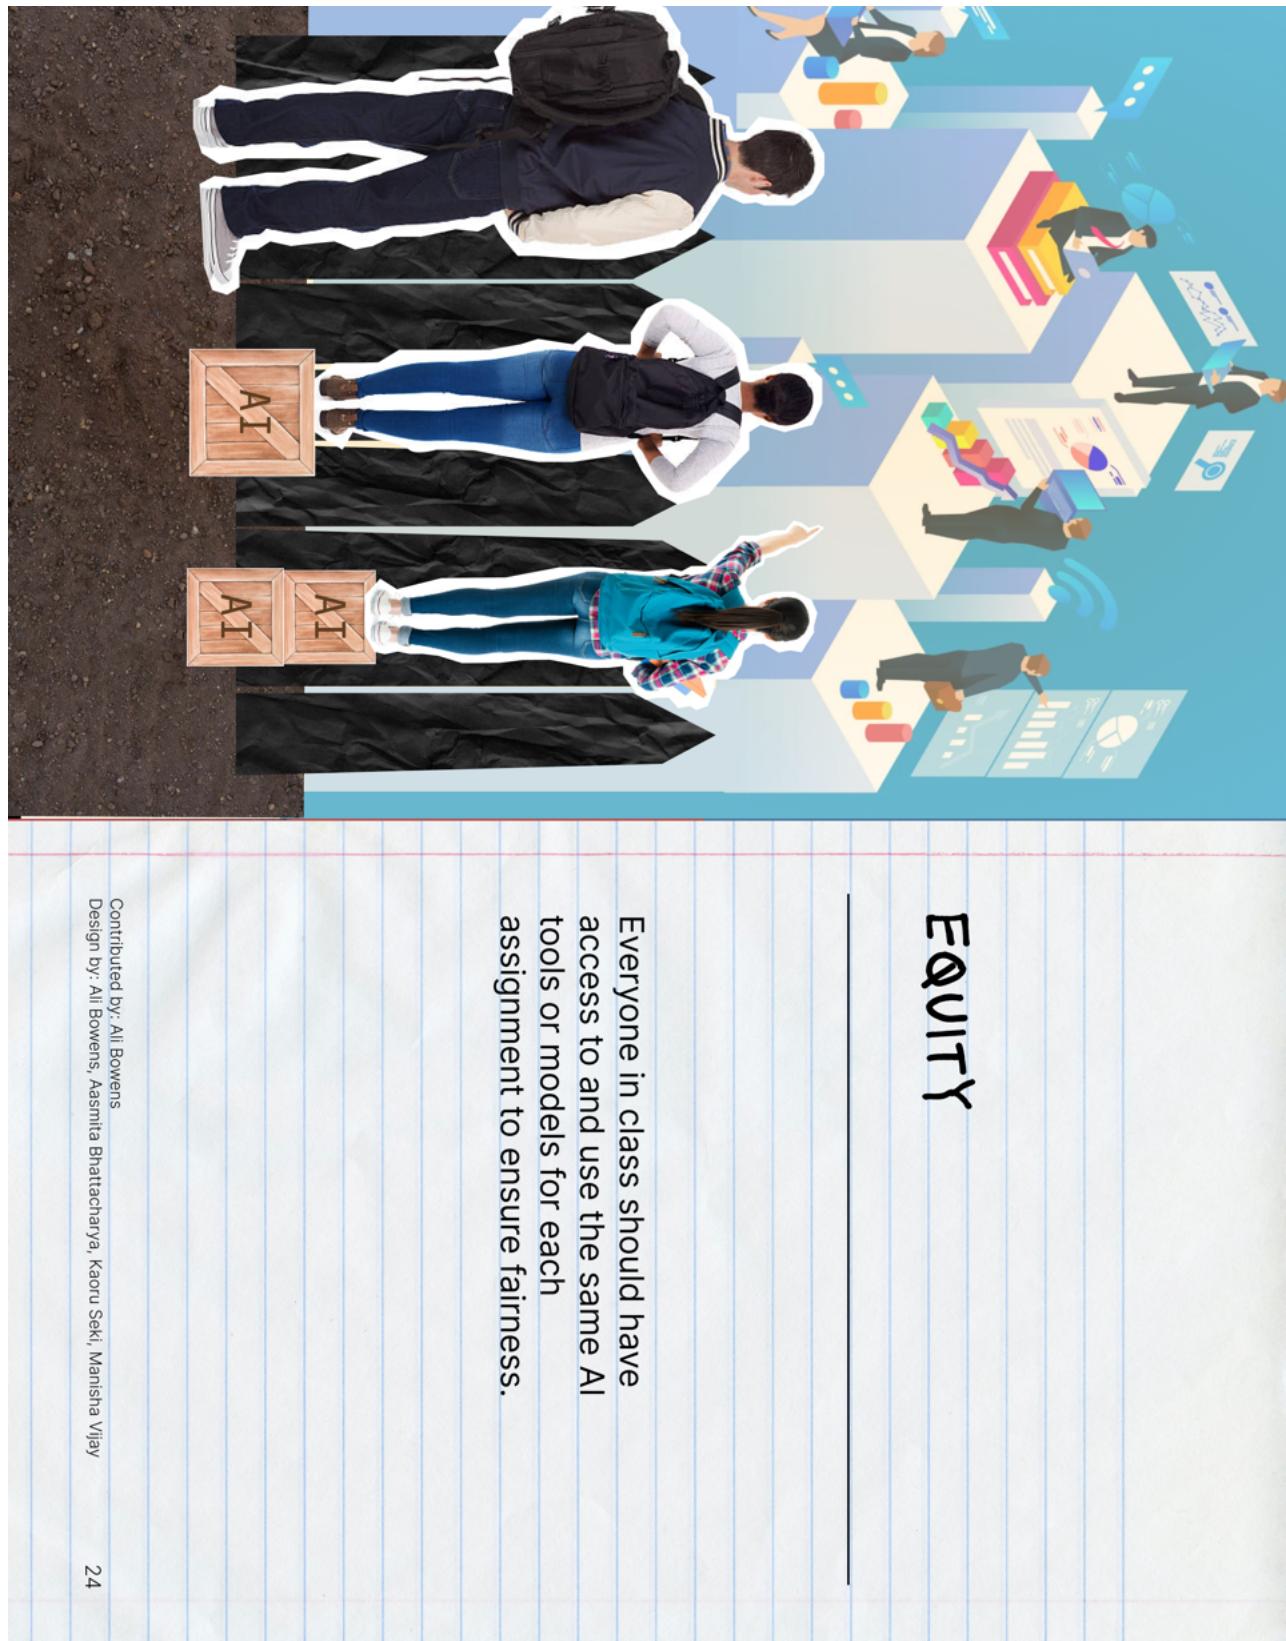


Figure A10: Equity, Page 10 of 10