Critical Engagement in Large-Scale Undergraduate Computing Programs

Jedidiah Tsang

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2024-137
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2024/EECS-2024-137.html

May 17, 2024

Critical Engagement and Belonging in Large-Scale Undergraduate Computing Programs

by Jedidiah Tsang

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination

Committee:

Assistant Teaching Professor Lisa Yan, Chair

Research Advisor

(Date)

17-May-2024

* * * * * * *

Associate Teaching Professor Joshua Hug

Second Reader

(Date)

17-May-2024

.

To my mentor, Lisa Yan: you have fundamentally changed the trajectory of my life. Thank you for your patience, your insights, and the way that you have cared about me and my well-being when teaching began to feel like a struggle. You have made this my best year at Cal.

To Joshua Hug: thank you for being my second reader (all the way in the Netherlands), and for teaching me that it's ok to be chaotic good sometimes.

To Su Min Park and Carol Li: you two have brought so much energy and unbridled joy into our research meetings. You both are talented, thoughtful, considerate individuals, and I can't wait to see what you do next.

To Zephyr Omaly, Ethan Ordentlich, SreeVidya Ganga, Allen Gu, Crystal Wang, Anton Zabreyko, Jero Wang, Eddy Byun, and Andrew Liu: co-instructing/co-head TA'ing alongside you has been the highlight of my time at Cal. Thank you for being such a source of inspiration and camaraderie.

To my family, teaching faculty, and course staff that have shaped my teaching and research journey in the most profound ways: thank you for your support and encouragement.

Table of Contents

Abstract

As the magnitude of the impacts that computing has on society exponentially increases, so has the interest in computing as a field of study. In this report, we explore the largest, open questions in computing ethics education by examining the most recent research in this area. We then establish the learning objectives and goals of the course, which fall largely in line with most computing ethics courses, with a particular focus on justice within the framing of our chosen readings. Beyond the learning objectives, we look at the ways that educators have chosen to approach pedagogy in this area. By setting the stage for *what* computing ethics education aims to achieve, and *how* we aim to achieve it, we can examine the usage of Question, Quote, Comment, or Response (QQCRs) as our chosen mechanism to reinforce course content. We do so by examining both course survey data and student QQCR responses from the Fall 2023 offering of CS195: Social Implications of Computer Technology. This leads us to the development of a codebook to classify themes that we observe among student work. To this end, we seek to measure how students critically engage with the course material, and how their perspectives on salient topics surrounding computing ethics shift across the term.

1 Introduction

Computing is ubiquitous—from cloud computing to machine learning, the federal government has invested billions of dollars annually on STEM education programs, and the California State Assembly is considering a bill to make Computer Science (CS) a high school graduation requirement in 2030 [12]. With all the buzz, the number of students majoring in a computing discipline has skyrocketed [38] across the U.S. Consequently, many CS educators have had to deal with the implications of teaching at scale. Especially at UC Berkeley, the number of students enrolled in CS2 Data Structures, one of our most popular lower division classes, has quadrupled since 2010 (Figure 1).

At the same time that computer science education has exploded in growth in recent years, the era of Big Data has driven substantial technological advancements. How computing technology has impacted society, and perhaps more crucially, how society shapes the technology that is developed and promoted, cannot be overlooked in computing curricula. As leaders in both industry and academia grapple with ethical considerations of computing technologies and roles these technologies have played in perpetuating and magnifying societal inequities [44], computing ethics curriculum has also become a core component of computing degree requirements. At every turn, the question of computing ethics, responsible computing [29], and social contexts are no longer purely a philosophical question; rather, these considerations are grounded in present-day reality.

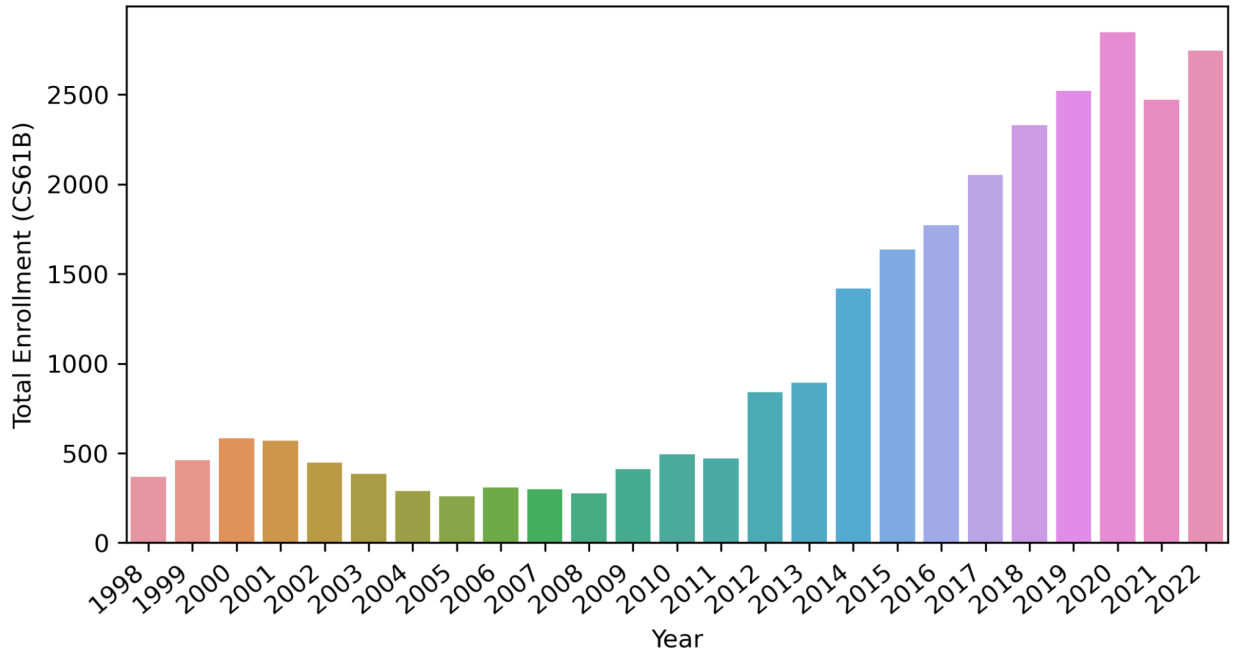

Figure 1. Enrollment numbers for UC Berkeley's data structures class over two decades

In this work, we define this area of computing education as *computing ethics*[1], where computing ethics curriculum often has two core learning outcomes. First, such curricula aim to *expose* students to a wide range of perspectives from a diverse population—ranging from subject matter experts to their peers—on issues that place the technology they create within the context of broader society. Second, computing ethics curricula also encourage students to *engage critically* with the material that the instructors have assigned so that they can formulate informed opinions on these issues.

With that said, teaching computing ethics classes poses a unique challenge compared to most other technical curricula; by nature, there are a vast array of topics that don't have a universally recognized "right" or "wrong" answer (e.g., AI regulation and safety). This sentiment is recognized by the original instructor of CS195, Professor Emeritus Brian Harvey: "[you] don't have to agree with me; what you have to do is show that you understand and take seriously points of view different from your own, and try to explain why your arguments are better than theirs" [16]. Many computing ethics courses have recognized the importance of discussion and collaboration as a means of successful engagement in the class [14], but this becomes especially difficult to achieve when higher education must scale.

For the CS2 course described in Figure 1, the course structure and curriculum have responded to the rapidly increasing demand due to the constraints on space, instructor time, and funding resources. Most adjustments pivoted to scalable strategies in an effort to support the most students possible. For instance, autograders were developed as a mechanism to deliver feedback to students without linearly increasing staff overhead, and office hours were limited to 10-minutes per request [18]. Consequently, these blanket policies run counter to the very foundation of a fair, equitable education—regardless of the studied benefits of an individualized tutoring approach [8], every student is now offered equal treatment, which cannot meet the variety of needs of their learning. On top of these existing barriers to equitable education, conventional tactics like developing autograders to holistically evaluate student work cannot be replicated nor translated easily in a large computing ethics course. Combining the scale that we teach at with the limitations of a one-unit course (translating to three hours of course engagement a week), it can be difficult to motivate students to make the necessary considerations to achieve our learning objectives.

These issues provide the impetus for the interventions we discuss in this report: A new curriculum for the computing ethics course at UC Berkeley (CS195: Social Implications of Computer Technology) engages students via assignments called Question, Quote, Comment, and Reply (QQCRs). After completing a set of assigned readings for the week, students are asked to either pose a question related to the material, name a quote they found significant (and why), make a comment about the reading, or critically, reply to an existing thread. By leveraging EdStem, a browser-based discussion forum, we allow students the ability to critically examine the material, share their lived experiences, and engage in discourse with their peers on topics ranging from generative artificial intelligence to memes and misinformation at scale.

To this end, we hope to address the following research questions in this thesis: **RQ1:** How does the structure of QQCR assignments impact students' critical engagement with the material? **RQ2**: How is a student's engagement with the material impacted by the particular topic

[1] Within this report, we'll frequently refer to computing ethics as a broader umbrella term, encompassing responsible computing, human context and ethics, and conventional ethics/philosophy.

being covered in a given week? **RQ3:** How does the framing of our course impact a student's beliefs on computing ethics considerations?

In this report, we explore the largest, open questions in computing ethics education by examining the most recent research in this area. We then establish the learning objectives and goals of the course, which fall largely in line with most computing ethics courses, with a particular focus on justice within the framing of our chosen readings. Beyond the learning objectives, we look at the ways that educators have chosen to approach pedagogy in this area. By setting the stage for *what* computing ethics education aims to achieve, and *how* we aim to achieve it, we can examine the usage of QQCRs as our chosen mechanism to reinforce course content. We do so by examining both course survey data (RQ3) and student QQCR responses (RQ1 and RQ2) from the Fall 2023 offering of this course.

2 Related Work

In this section, we compare CS195 to existing curricula and course structures in three ways. First, we examine how CS195 fills the role of teaching “computing ethics" in a student’s broader computing curriculum. Next, we assess the specific assignments used in the course and the limitations of different approaches, leading us to propose QQCRs as an assignment archetype that fills a gap identified by previous researchers. Finally, we critically examine the learning objectives of our computing ethics classes in order to determine the methodology by which we measure the efficacy of the class on such learning outcomes.

2.1 Computing Ethics Course Curricula

Recent advancements surrounding teaching computing ethics in higher education computing curriculum have proposed one of two strategies: integrate “embedded ethics” modules into a student’s technical coursework [7, 22], or invite educators from various disciplines in a standalone course [42]. One of the common objections to both strategies is a concern of pragmatism: that teaching computing ethics takes time away from teaching technical topics [27]. Regardless of the basis of this dominant narrative, the Electrical Engineering and Computer Sciences undergraduate curriculum at UC Berkeley (EECS) has chosen the latter approach by designating an Ethics Requirement [21], where students “complete one course about engineering ethics or the social implications of technology…for a ‘P’ grade.” for undergraduates studying Electrical Engineering and Computer Science (EECS). Courses range from home departments such as Bioengineering and Science, Technology, & Society to Public Policy.

To satisfy the Ethics Requirement, most EECS students choose to take CS195: Social Implications of Computer Technology, which minimizes the requisite work and offers the most flexibility as a 1 unit, pass/fail course. In other words, students are only expected to spend 3 hours a week (including lectures and deliverables) on the class, and can only take it for a pass/fail grade (as opposed to letter grade). This grading decision is intentional—to allow students the ability to have frank, candid dialogue without the penalties of having the "correct" opinion. From the course webpage, "the reason for that policy is to ensure that you can feel free to express opinions different from those of the instructors, both in class meetings and in written work." Optionally, students can enroll in CSH195, which is a 3 unit graded course. Both courses were developed by Brian Harvey (Teaching Professor Emeritus) at UC Berkeley. Given the structural limitations of this course, in our work much of our interventions focus on prefiguration [9], or assessing what students know about a given computing ethics topic. With the general framework of this class established, we can move into the modes of participation in the course.

2.2 Classroom Engagement and Participation

A survey of the past four decades of standalone computing ethics courses divides pedagogical activities into three main buckets: participation (consisting of discussions and active learning), knowledge exposure (lectures and reading), and submitted work (via assignments and projects) [39]. Approximately 44% of all courses within this survey incorporate all three activities within their course work [39]**.** Within the context of participation and exposure, classroom interactions in a *synchronous* format tend to benefit from exposure of different

perspectives among students. However, such synchronous work is inherently constrained; within lecture, “fifty students [are] too many for general discussion” [11] but small group discussions inherently limit the amount of perspectives that students are able to engage with.

Consequently, alternate ways of participation in an *asynchronous* format include discussion course forums. Course forums are an effective way to continue classroom dialogue beyond lecture [46] and expose students to other perspectives, increasing self-efficacy and critical thinking of the topics at hand [29, 47]. With a class size of several hundred students, course forums can operate at scale by leveraging the diversity of background in support of computing education learning objectives. Combined with readings (which are conventionally passive activities), we can empower students to utilize course forums as a method of accountability for exposure (learning new concepts and theories), as well as turning reflections on course forums into a form of submitted work. We’d like to emphasize the significance of using course forums as a method of exposure, since it can directly translate into student learning outcomes. When thinking about exposure, we can bifurcate this concept into both exposure to expert opinions and the perspectives of their peers. With the latter in particular, the concept of social annotation provides the impetus for the pedagogical value of having peer to peer interaction. Social interaction is defined as “a genre of learning technology that enables the annotation of digital resources for information sharing, social interaction and knowledge production” [32]. Asking students to post asynchronously on course forums *and* respond to one another meets the definition of social annotations, and can provide a student with greater insights into the topic they’re learning.

2.3 Equity, Belonging, and Justice in Education

With any computing ethics course, how the instructors choose to frame the course (ranging from the selection of topics, to the granularity of individual readings) holds important ramifications on student engagement and outcomes. Recent literature in computing education has focused on the concept of a justice-centered approach to education, and how our pedagogy can draw from culturally relevant material to expose students to their work in a larger societal context [34]. The Research in Equity and Sustained Participation in Engineering, Computing, and Technology (RESPECT) conference was formed specifically to focus on how we can integrate these values into our teaching. As a result, in the research process, we sought to incorporate aspects of a justice-centered education in both the course structure, as well as the methodology in our qualitative analysis. To this end, we focused on two concepts: cultural competence and critical consciousness.

Cultural competence is defined as “a multifaceted concept that encompasses the ability…to effectively understand, respect, and respond to the diverse cultural backgrounds and identities of students, families, and communities…” [14]. Building cultural competence is particularly relevant in a computing ethics course in a globalized society, where our students will undoubtedly be working alongside and serving populations from various cultural backgrounds. However, even in the context of the classroom, building cultural competence allows educators to support students’ diverse methods of learning and engagement. Especially in a computing ethics course where discourse frequently yields civil disagreement, integrating cultural competence and measuring the ways that students interact with different cultural backgrounds allows us to reach our learning objectives. This can be operationalized via the selection of readings to ensure that

students are exposed to a wide range of cultural perspectives, and the methods of interaction as well (encouraging group work, discussions, and multimedia resources [14]. Cultural competency as a phrase is used interchangeably with cultural humility; we note that while the former places the onus on the individual to learn, the latter treats cultural knowledge as a dialogue and learning process.

Although prioritizing cultural competence in the context of our course has value in and of itself, building cultural competence also allows students to think critically about systems of power that marginalize certain populations. This idea is captured by Paul Freire's concept of *critical consciousness*, which encourages us to consider how education can be used as a catalyst for recognizing and taking action on systems that marginalize groups [40]. This tool is particularly useful for members of historically underrepresented groups to recognize systems of injustice, and connect that to better student achievements and outcomes [20]. Building critical consciousness also matters for people in positions of power and authority, as they are better able to serve a diverse population through the work they do [6].

Within the context of a computing ethics course, building critical consciousness involves exposure to material beyond just looking at different cultural backgrounds; it also includes material on race, gender, and other affinity groups that can result in disparate outcomes. To this end, we can thus both support students within their own academic journey (by making their stories feel heard and validated), but also encourage all students (not just minoritized students) to think critically about how the technologies that they create can directly impact social systems that perpetuate inequities. As our students go on to build technologies in Silicon Valley and for the government, it is prudent to educate the majority on how to both recognize and take action on those aforementioned systems.

3 Course Structure

CS195: Social Implications of Computing Technology has two primary student learning objectives:

(1) Gain exposure to contemporary topics at the intersection of computing and society that are directly applicable to student experiences during and after their studies; and

(2) Engage critically with and form opinions on course topics via discussion with peers. Because course topics are presented primarily through frameworks of critical consciousness and justice, a secondary learning objective is therefore to teach cultural competence (also known as cultural humility; see Related Work).

The full course syllabus is in Appendix I; this section highlights the salient portions of the course structure that lend to the learning outcomes we hope to achieve, as well as the limitations of our work.

3.1 Course Activities

The course is 1 unit, pass/fail, corresponding to about three hours of work a week. There are three primary course activities (Figure 2): one 80-minute weekly lecture, weekly reading assignments—Question, Quote, Comment, and Reply (QQCRs), discussed in the next subsection—and two short, paired podcast assignments. Surveys were administered weekly to provide space for students to consider their stance on that week's course discussion; while most were anonymous for the purposes of starting lecture discussions, three longer surveys were required for a student's passing grade.

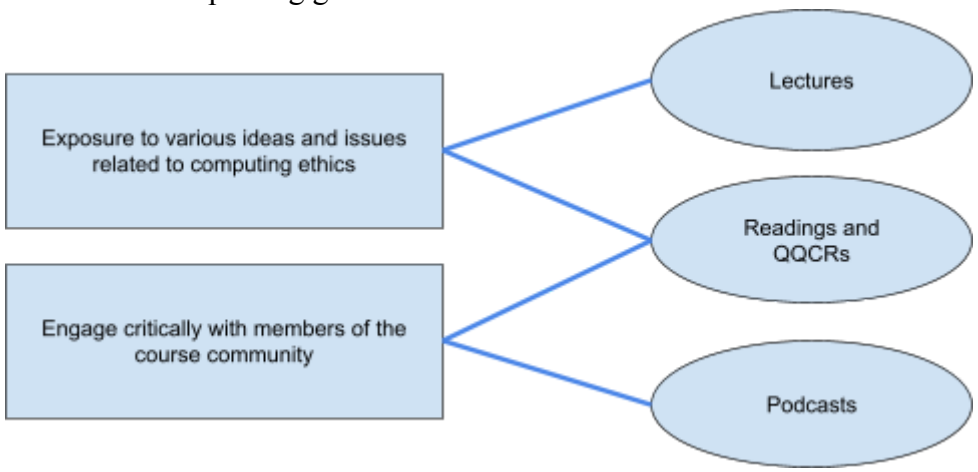

Figure 2. Course objectives mapped to course activities

3.2 Content Schedule

Each semester's course topics approximately follow the ordering as listed in Table 1, where one topic is presented per weekly lecture. The course first covers course topics common to similar courses in other institutions [23] before explicitly introducing social frameworks that define much of the "tech industry." A three-week artificial intelligence (AI) unit was developed for Fall 2023 to provide space to parallel contemporary discussions on generative AI, a technology that started making news headlines in late 2022. The semester ends with a "grab bag" set of topics that are often extensions of earlier topics. Lectures were delivered primarily by the course instructor in Electrical Engineering and Computer Sciences (e.g., in Fall 2023, an Assistant Teaching Professor). Due to the scale of the course and physical classroom capacity, lectures were delivered hybrid (in-person lecture and simulcast online). Lecture attendance was required, and anecdotally about 15% of students opted to come in person.

Technology and Society	Week 1: Free Speech Week 2: Social Connection and Time Well-Spent Week 3: Data Privacy Week 4: Software Risks and Professional Ethics
Social Frameworks of the Tech Industry	Week 5: Technosolutionism Week 6: Government Censorship and Surveillance Week 10: Silicon Valley and the US Week 11*: Tech and Labor
Artificial Intelligence	Week 7: The AI Debate Week 8: Google AI Ethics Week 9*: Generative AI, Part II
Topic Extensions	Week 12*: CS Education Week 13: Memes, Misinformation, and Media Literacy Week 14: Tech and Labor Part II, Closing

Table 1: Course topic schedule by week. Weeks denoted by an asterisk (*) had lectures delivered by guest speakers with domain-specific expertise.

3.3 QQCRs

Given the short amount of class time in the course, the bulk of student activities were focused on weekly readings via the Question, Quote, Comment, and Reply assignments (QQCRs, for short). Every week has 1-3 readings associated with that topic, and prior to class students are expected to post on a public course forum (although they could post anonymously if they wished to). Each week, each student must post a minimum of 1 reply to another student's QQCR, and 2 other QQCRs. Notably, this requirement was adjusted—in early weeks of Fall 2023, students submitted up to 5 QQCRs, which increased workload. Students could consolidate all 3 of their QQCRs in one reading for that week, or distribute them across multiple readings. All QQCRs should be completed prior to the lecture itself, although the Fall 2023 assignment lateness policy allowed for students to submit all their QQCRs on the last day. When grading the QQCRs, staff were extremely lenient - other than exceptions for one word responses, all QQCRs received credit. There is a stark difference between the amount of effort put into QQCRs disaggregated purely on submission time; we discuss this more in Section 5.

3.4 Community Guidelines

A public course forum necessitates community norms; for CS195 these were described in the syllabus as below:

> Our classroom environment extends to online spaces like Ed. Be civil and respect others, even those that are not in this course. There are 3 ways for you to address hateful comments and prevent harassment, in order of decreasing emotional labor (the last two do not count towards QQCR): reply on Ed directly to post; submit an anonymous report form to course staff; or reply on EdStem with "That's not cool."
>
> Course staff reserve the right to delete and/or not award credit or dismiss students from seminar for remarks that promote hate speech, trolling, harassment, discrimination, or defamation. Serious offenders will be reported directly to both the CS Department Grievances faculty member and the University Center for Student Conduct.

3.5 Other Submitted Work - Surveys

Students were asked to submit three course surveys in Weeks 1, 6, and and 15. Surveys were not anonymous, although we de-identified the data for the purposes of our report.

3.6 Other Submitted Work - Podcasts

Students also submitted two 7-10 minute podcasts across the semester. Through the podcasts, students were able to synchronously engage with a peer on a topic of their choosing (relevant to the course material), and peer review the podcasts of other students as well (which served as exposure and decreased staff overhead). Podcasts served as a replacement for essays, as "most of you will engage with these implications through spoken conversation, we're hoping that this will be a better simulation of what you will do in practice—as opposed to a written essay" [17].

4 Methodology

4.1 QQCR Codebook Creation

Our QQCR data came from scraped EdStem data from the CS195 Fall 2023 offering, removing all instances of Personally Identifiable Information (PII) to de-identify our dataset before performing our analysis.

To capture the various themes related to our course learning objectives, we had to consider how various factors of a QQCR signaled engagement (critical or not) with the material. Although our codebook was developed after we had collected the data, our method largely followed the grounded theory approach [28] as follows: First we examined the data independently to generate an initial iteration of the common themes. We then developed theories on the causal mechanisms resulting in the themes we observed. Finally, we iterated on the initial codebook to reach a consensus on the coverage to ensure that they aligned with the course concepts and learning objectives.

Our codebook as listed in Figure 3 identifies the following themes that we focused on within the class: Ethical Framework, Judgment, Attitude, Critical Reasoning/Rationale, Cultural Competency Continuum, Positionality, Critical Consciousness, and Social Systems. A majority of these themes are constructed through course learning goals and related work as we describe next.

In alignment with the focus on a justice-centered education, a portion of our codebook centered on two themes: Cultural Competency Continuum [19] and Critical Consciousness (see Chapter 2: Related Work) [36]. We removed levels of granularity that seemed difficult to parse within the context of the QQCRs, which typically ranged from 60-100 words each. These two themes were especially prevalent in Week 5: Technosolutionism (with the reading Race After Technology [43]), Week 6 (Government Censorship and Surveillance), Week 10 (Silicon Valley and the U.S.), and Week 13 (Memes, Misinformation, and Media Literacy). Given that many of the students were offering personal anecdotes or making references to their own identities, we added an additional theme on Positionality; in other words, how a given student's identity (race, gender, sexual orientation, etc) affects their standing with society (or additionally, the intersections of multiple identities) and consequently, the positionality of their QQCR response.

Beyond this, given that one of our learning objectives is the *critical engagement* with course material, we observed students frequently express some normative opinion of the world around them (or of themselves). We created the Social Systems theme to see what facet(s) of society students choose to talk about. Also related is the Judgment theme, with a focus on whether the student also identifies themselves as a subject of the judgment. We also sought to see *how* students defended their claims in addition to the claim itself (Critical Reasoning/Rationale), as well as whether students were pessimistic or optimistic about future prospects (Attitude).

One final theme we identified was the nature of the student's QQCR with regards to a potential underlying ethical framework. In this sense, we were able to record data on a student's judgment beyond *who* the opinion is expressed on, but also *what* the opinion is about. To this end, we created two subthemes to represent a justice-centered lens and utilitarian-centered lens (which are diametrically opposed to one another), and certain lenses that capture the spirit of human collaboration/humanity (common good/virtue lens), as well as one focused on the rights

of individuals (rights lens). We drew from Santa Clara University's framework for ethical decision making [48], although we only applied the lenses we saw most frequently.

Subthemes	Code	Description
Ethical Framework		
The Rights Lens	ER-R	Ethical actions protect the inherent moral rights and dignity of all involved.
The Justice Lens	ER-J	Ensuring that individuals receive what they are due, be it in terms of social structures, distribution of resources, or rectifying wrongs.
The Utilitarian Lens	ER-U	Evaluates the ethicality of an action based on its outcomes, aiming to maximize the good for the majority.
The Common Good / Virtue Lens	ER-CV	Promotes ethical actions that contribute to the communal well-being, emphasizing societal interdependence.
Judgment		
Individual Judgment	J-I	Judgment about oneself OR another individual.
Interrelational Judgment	J-IR	Judgment about oneself AND other individuals/groups.
Broader Social Judgment	J-S	Judgment about society in general (the 'public') or groups of individuals without direct reference to self.
Attitude		
Optimism	A-O	The student addresses a potential solution.
Pessimism	A-P	The student explicitly states that there are no viable solutions or paths forward.
Critical Reasoning/Rationale		
Pathos Appeal	CR-P	Using personal anecdotes or other appeals to emotion to back their reasoning.
Ethos Appeal	CR-E	Using an appeal to authority to back their reasoning.
Human Nature/Philosophy	CR-H	Backs up their judgment by making some assumption/claim about human nature.
Logos Appeal	CR-L	Using statistics/figures or logical inferences to back their reasoning.
Cultural Competency Continuum		
Cultural Homogeneity	C-H	Opposed of cultural competence - believes in ethnocentrism and holds that a given culture should be widely adopted in other spaces.
Cultural Blindness	C-B	Holding the belief that racial or cultural differences should make no difference.
Cultural Competence	C	Agencies or systems that acknowledge and respect cultural differences.

Figure 3. Codebook used for QQCR Analysis

Subthemes	Code	Description
Positionality		
Vacuum	P-V	The student understands one's place in society, but not how their identities may influence that place.
Affinity	P-A	The student understands how a given affinity group (race, gender, ethnicity) affects one's positioning in society.
Intersectionality	P-I	The student understands how multiple affinities interact with one another to affect one's positioning in society.
Critical Consciousness		
Critical Reflection	CC-R	"Process of learning to question social arrangements and structures that marginalize groups of people" [36].
Critical Motivation (Political Efficacy)	CC-M	"Perceived capacity and commitment to address perceived injustices " [36].
Critical Action	CC-A	"Engaging individually or collectively to change perceived injustices" [36].
Social Systems		
Family and Friends	SS-FAM	Accounting for kinship ties and close relationships.
Tech Industry	SS-SV	Mentions the structure of how the private tech industry operates.
Education System	SS-ED	How education works (what we teach and how we teach it).
Economic System	SS-EC	Contains points about markets, corporations, and distribution of goods/services.
Political System	SS-POL	Talks about laws, regulations, and government processes.
Other		
Non-US Comment	NI	Discusses a non Western perspective in their comment.
Summary	S	Either **exclusively** summarizes the reading OR (if reply) **exclusively** agrees with a parent comment.
Inconsistent	I	Inconsistent arguments/claims/logic throughout their discussion.

Figure 3, cont. Codebook used for QQCR Analysis

4.2 Coding Process

After finalizing our codebook, we first hand-coded a few hundred responses. To obtain standardization across various codes, we agreed on various thresholds to determine when a code was appropriate for a given QQCR. We recognize that many of these codes can be subjective and so there was some inevitable variation in how exactly the codes were applied. Below is an example of a set of codes we applied to a QQCR:

> ***Student Post***: [Quote] " 'When images of Black people being killed by police garner over 2.4 million clicks in 24 hours, and the average "cost per click" for related content reaches $6 per click [Image 1.3], the virality of Black death is not only incentivized, but nearly guaranteed.' "
>
> The naivety in me forgets that this is all for money and that somehow makes it even worse. I understand technology is profit-driven and these private companies do things for money, but I'm wondering what steps we can take to encourage legislators to actually enforce that search and social media algorithms at least provide censorship to all people. I think there's also something that needs to be done in our education system to stop this perpetuation of normalizing Black suffering.
>
> ***Codes Applied by Hand***: ER-R (Rights Lens), J-S (Societal Judgment), A-O (Optimism), CR-P (Appeal to Emotion), C (Cultural Competence), CC-R (Critical Reflection), SS-SV (Tech Industry), SS-SV (Education System), SS-EC (Economic System), SS-POL (Political System)

After hand-coding a few hundred responses, we could then train a model to extrapolate our hand-coded responses onto the remainder of the data (we had approximately 9000 QQCRs across the Fall 2023 term). To determine which model we would use, we fine tuned Open AI's `gpt-3.5-turbo` model on 156 pieces of hand-coded QQCRs, then manually compared the outputs on a few QQCRs to Open AI's `gpt-4` model, as well as Hugging Face's default facebook/bart-large-mnli model. Both the `gpt-4` model and the `bart-large-mnli` model were used as zero-shot classifiers, but all three models received the same context. Notably, the context we provided only contained the "Subthemes" column due to the sheer length of the whole codebook, as well as some formatting instructions. Given that the gpt-3.5-turbo model was fine tuned, we saw much less variance in the outputs of the model across multiple requests, and performed much better with the context provided. If we provided the full context (the entire codebook including full length descriptions), we found that the `gpt-4` model performed the best. For future work (Chapter 7), we would like to explore some medium between just providing the name of the code and providing the whole codebook, and more rigorous analysis on the exact model to use/fine tune for this task.

Although this approach allows us to classify large amounts of QQCRs, this also limits the types of research questions we can answer using the extrapolated data. Across the 8,545 entries that were classified using the model (after filtering for students who opted out and staff comments), the `gpt-3.5-turbo` model hallucinated codes that didn't exist ~60 times, or

around 0.02% of the outputs. Overall, the exact codes applied with the fine-tuned model differ slightly compared to hand-coder expectations. For example, below is a QQCR that was classified by the model:

> ***Student Post***: "[Reply]: I think this really highlights the idea of how mainstream AI is bound to become in our lives within the next decade. As you already mentioned, ChatGPT has the ability to ""innovate"" only by manipulating data that it has already learned upon with some measure of ground truth information. I'm more excited than scared to see what the future of AI holds for us civilians. The reasoning behind that is because I feel that our government is likely a decade or two ahead of us in technological advancements, so if AI is to be used for militarization and racial profiling, it's probably already happening right under our noses, and we've become rather accustomed and complacent to it."
>
> ***Codes Applied by Model***: ER-J (Justice Lens), ER-U (Utilitarian Lens), J-S (Broader Social Judgment), CR-H (Human Nature/Philosophy), CR-L (Logos Appeal), SS-POL (Political System)
> ***Codes Applied by Hand***: ER-J (Justice Lens), J-IR (Interrelational Judgment), A-O (Optimism), CR-H (Human Nature/Philosophy), SS-POL (Political System)

Given the above, any sort of analysis that involves looking at the *exact codes* applied to a reading were not appropriate with our model. Instead, we treated *the number of codes* the model outputted as a proxy for the degree of student engagement.

4.3 Codebook Creation for Assigned Readings

One other application of the codebook was to apply them onto the *readings* themselves to see if there was a trend between the codes that were applied to a given reading, and the ways that students then engaged with that material. The coding process was similar to the process of coding student QQCRs, with one notable exception: Do we apply codes only to the author's direct claims and observations, or to all sides presented? Given our focus on exposure as a learning objective, we chose the latter. One example would be current debate around a pause on AI advancement [45], which presented both arguments, thus increasing the number of codes applied to that reading. A comprehensive list of all the codes that were applied to course readings can be found in Appendix III.

4.4 Survey Data Analysis

Research involving the survey data utilized a longitudinal survey design, collecting data from all students enrolled in CS 195 for the first and final weeks of the Fall 2023 semester. This survey consisted of 27 five-point Likert-scale questions ranging from "Strongly Disagree" to "Strongly Agree". Topics were grouped into four sections, and asked students their attitudes towards significant claims related to the course content. The Likert scales used in the first and final weeks were deliberately designed by the instructor to differ; the final week's scale excluded the "Neutral" option to encourage students to make definitive judgments. The de-identification

process for survey data followed similar procedures to the QQCR process: we removed all PII, and shuffled the de-identified data controlling for two demographic factors: graduation year and gender. For example, all responses from a female student graduating in 2024 were swapped with another female student graduating in 2024. In order to perform cross tabulation, we ensured that the same students were swapped in both Week 1 and Week 15.

The design of the survey questions align with relevant topics covered throughout the semester, and framed within the context of the ACM Code of Ethics [1], which emphasize the importance of computing professionals' actions in supporting the public good and minimizing harm. It further highlights CS195's commitment to integrating computing ethics into the curriculum and examines whether exposure to these ethical discussions enhances students' critical thinking skills and awareness of social issues [15].

Data analysis was performed with Python libraries such as Pandas for data manipulation, and Matplotlib, Seaborn, and `plot_likert` for visualization. The analytical techniques employed in this study comprised diverging bar stacks and cross-tabulation, the latter of which was specifically applied to three key themes discussed in the course: ethical framework, critical consciousness, and technology and labor. This approach facilitated a detailed visualization of shifts in student responses, particularly highlighting the transition from the "Neutral" option to other available choices in the Week 15 survey.

5 QQCR Analysis

With the QQCRs, we aimed to answer **RQ1**: How does the structure of QQCR assignments impact students' critical engagement with the material? and **RQ2**: How is a student's engagement with the material impacted by the particular topic being covered in a given week?

5.1 Interactions with Integrated Flexibility on Student Engagement

One policy decision with QQCRs is the ability for a student to choose which readings they distribute their weekly QQCRs across. In other words, in a given week, a student could choose to distribute all three contributions equally, or post three times on the discussion thread for one reading.

We do further analysis on Week 9, which had a high amount of variance in the codes that were applied to the reading:

Science, "Generative AI meets copyright" [41]. 2023.
Codes Applied: ER-R, ER-U, J-S, CR-L, SS-SV, SS-POL

The New York Times, "35 Ways Real People Are Using AI Right Now" [25]. 2023.
Codes Applied: SS-FAM, SS-ED

Guardian. "AI machines aren't 'hallucinating'. But their makers are" [37]. 2023.
Codes Applied: ER-R, ER-J, J-S, CR-P, CR-E, CR-L, SS-SV, SS-EC, SS-POL

Taking a look at Week 9, we plotted the distribution of students across three readings, as shown in Figure 4.

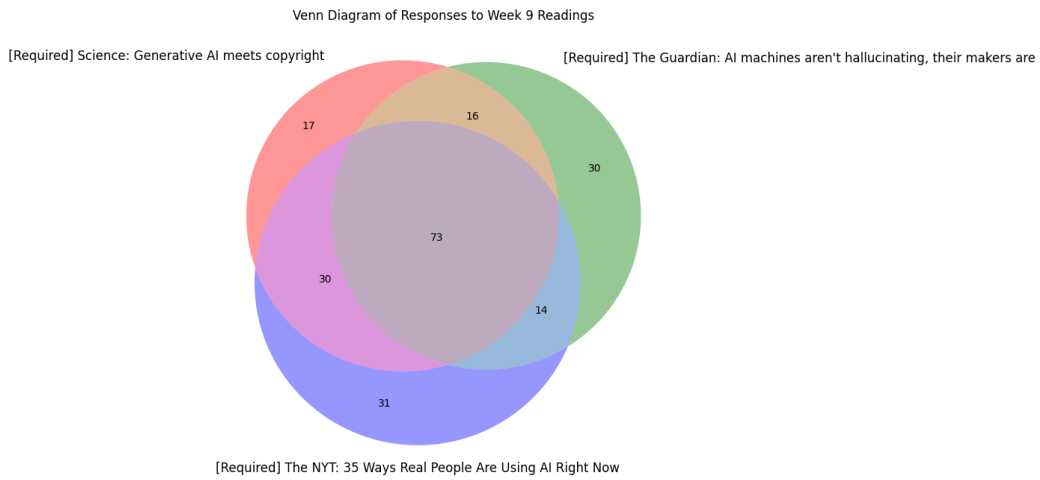

Figure 4: The distribution of comments across the three Week 9 readings. Each circle represents one reading; overlapping regions indicate that the students posted a QQCR on those readings.

Out of the 211 students who submitted a QQCR in Week 9, 34.5% of them (73 students) posted a QQCR on all three readings, and 37.9% (80 students) posted on two readings. While we admit that students may be leveraging the course policies to reduce their workload, this reading completion seems to suggest that students are achieving (in part) one of our course objectives, which is to expose students to various perspectives. Additionally, at first glance, it may be concerning that 31 students chose to only respond to *The New York Times* Article "35 Ways Real People Are Using AI Right Now." Because this article has the least codes applied across all three readings, it may be natural to conclude that 31 students lacked critical engagement with course material. We discuss in the next section why this may not be the case.

5.2 Engagement Based on Topic

In this section, we use the number of codes applied to QQCR responses for a particular reading to assess students' levels of critical engagement with the material.

Among the readings with the least codes applied to student QQCRs (Table 2), the commonality is that most of these readings do not provide normative judgments on how the world should be (Case Study 1, Pew Research, Logic Magazine, and Carolyn Chen). Even with hot topics like generative AI, it seems like students generally aren't engaging deeply with the readings that don't have strong opinions themselves. One other note is that many of the readings with the least codes (contrasted with the readings with the most codes) don't have personal anecdotes attached to many of these stories.

Another noteworthy result was that the reading on the making of the tech worker movement had 13 codes (Appendix II) assigned to the reading (the third highest out of all the readings), despite having one of the fewest codes applied to the QQCRs. This reading was assigned in the last week of the semester, and was one of the longest (over 12,000 words long). It is highly unlikely that students read the entire text (and thus engage with every aspect of the text). Further examination on the relationship between reading length and the number of codes applied could yield some significant correlation as well.

The most interesting result was the fact that one of the readings with the fewest codes assigned to it ("35 Real Ways People Are Using AI Right Now" [25]) was in the top five readings with the *most* codes applied to the QQCRs associated with that reading. This seems unintuitive at first glance and perhaps suggests a lack of critical engagement:if an author does not go into depth in its own critical analysis of the topic, why do the same? A partial explanation can be that students are generally more interested in contemporary media, and issues that are relevant to them generally result in genuine engagement with the material [5]. However, this doesn't completely explain why students are engaging *deeply* with the material; beyond just what the codes suggest, much of the dialogue on the course forum also went far beyond the surface level usages of AI. We postulate that this can be attributed to a few various factors: first, this reading was assigned in Week 9. By this point, students had already had two weeks on Generative AI (Weeks 7–8), and may have developed opinions that they did not previously hold. Second, this reading was *not* dense; ironically, the fact that this reading did not have deep insights allowed a greater variety of students to engage with the material by sharing personal experiences and anecdotes. This dispels the notion that readings necessarily need to have dense, heavy analysis associated with them in order to encourage critical engagement with the material

(although it is still relevant to assign these readings to achieve exposure, one of our learning objectives).

Reading Title	**Average Number of Codes Applied to Student QQCRs**	**Number of Codes Applied to Reading**
Case Study 1: Use of AI in education in a research context [2]	2.67	3
Nature: How social media affects teen mental health [4]	2.74	3
Pew Research: Teens and social media [24]	2.78	1
Logic Magazine: The Making of the Tech Worker Movement [10]	2.8	13
Carolyn Chen: Work Pray Code, Introduction [13]	2.83	6

Table 2: The five readings with the least average number of codes applied to QQCRs

Reading Title	**Average Number of Codes Applied to Student QQCRs**	**Number of Codes Applied to Reading**
Nylon: For plus-size creators, Tiktok presents a new wave of challenges [26]	4.81	11
Science: Generative AI meets copyright [41]	3.49	6
The New York Times: The Rise of the Worker Productivity Score [31]	3.49	3
The New York Times: 35 Ways Real People Are Using AI Right Now [25]	3.38	2
Slate: The Cruel New Era of Data-Driven Deportation [3]	3.32	8

Table 3: The five readings with the most average number of codes applied to QQCRs

5.3 A Big-Picture Snapshot: Student Engagement over Time

Figure 5 portrays our metric of student engagement: the average length of responses from all QQCRs for each week in the semester. We decided to use average length as a proxy here, since longer responses are typically a signal of higher engagement or interest in the material. Given the duration of the semester, we would expect student engagement to taper off as the semester progressed, so any outliers were of particular interest to us. Thus, we sought to explain the "peaks" of student responses in Weeks 7, 9, and 11. We note there was a drastic increase in student engagement from Weeks 2 to 5, the change in QQCR completion policy early on in the term may have affected this trend. Because the syllabus changed to require less QQCRs starting Week 2 (Section 3.3), it may have allowed them to be more reflective in the posts that they did write.

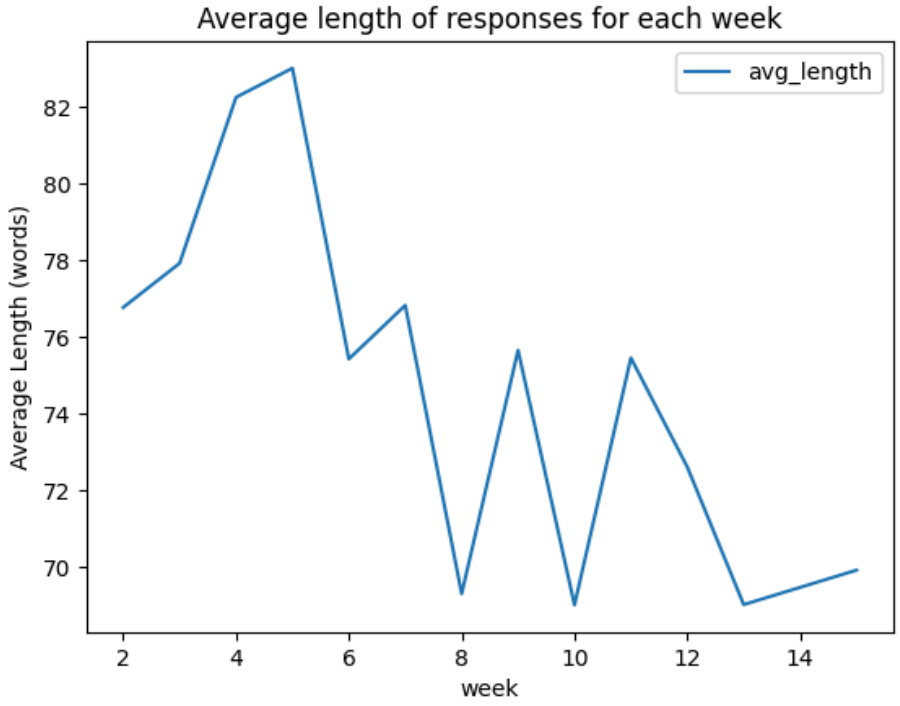

Figure 5: Average length of student QQCRs across the Fall 2023 term

Weeks 9 and 11 had higher engagement, and context in regards to current events may be helpful to understand this trend. In Fall 2022, the United Auto Workers (UAW) academic workers underwent a strike for higher annual base salaries for postdoctoral scholars, academic researchers, academic student employees and graduate student researchers. Many professors adjusted their methods of instruction to stand in solidarity with the strike, which looked like releasing pre-recorded content, changing assignments to optional, and modifying exams to

different formats as well. This change may have affected students in CS195 by realizing that the strike directly impacts their education, which allows them to be more engaged in this event. This is also reflected in Cowan's work, which states that "We asked students to examine the readings for each topic and tell us their preferences…The topics chosen were concrete and in the news, such as social media and career…" [11]. Cowan's students inclination toward relevant, current events can also explain the higher engagement response for students in CS195 during the strike.

It is helpful to understand any important trends or factors that lead to low response rates for certain weeks. Specifically, Lecture 8 covered Google AI Ethics, but the lecture was cancelled. Students were still expected to complete their QQCR assignment, so a synchronous learning week may have affected student performance, whether it may be digesting reading content or not having been exposed to the subject matter multiple times before writing their QQCR's. Not having a lecture for Week 8 may have caused students to be less engaged in the material, as there was not a discussion that was demonstrated by an instructor lecture. However, one important point Cowan brings up is that having instructor-led lectures may persuade students to view social implications topics in a particular manner: "Students speaking in class inevitably seek the instructor's approval, unconsciously assessing the instructor's bias" [11]. It may be worthwhile to assess whether there is less of a consensus between students on weeks where there is no lecture, versus when there are.

We provide a quick aside to discuss the impact of coursewide extension policies on student engagement. As aforementioned, students were expected to complete the QQCRs for a given week by the day of the lecture covering that week's material. However, students were allowed to submit work late (up until the end of the instructional period) without penalty. Disaggregating the above analysis (Figure 6) allows further insight into the stark difference between length of submission for submissions before and after the deadline.

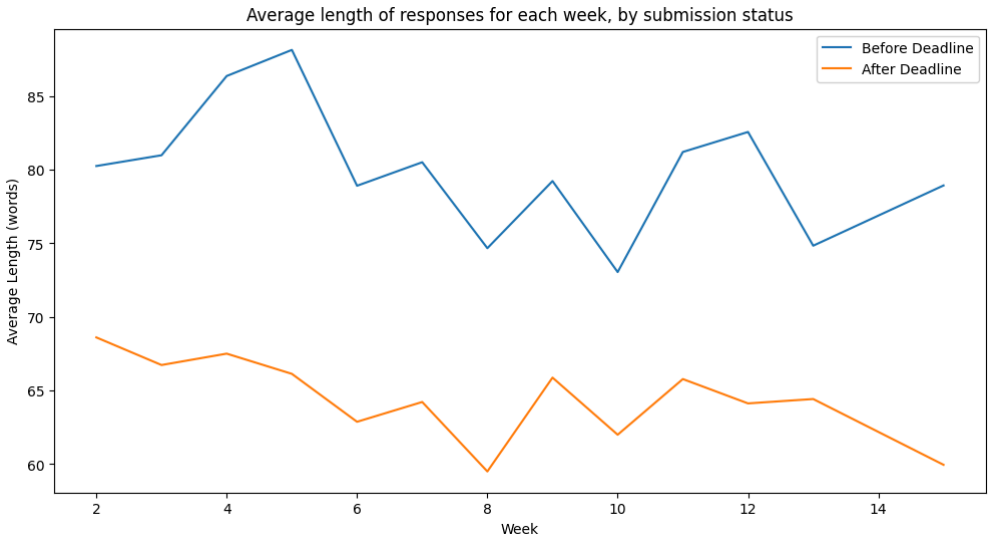

Figure 6: Average length of student QQCRs across the Fall 2023 term, disaggregated between students who submitted before the day of the lecture, and those who submitted afterwards.

6 Survey Analysis

With the survey data analysis, we aimed to answer **RQ3**: How does the framing of our course impact a student's beliefs on computing ethics considerations?. The analysis of student responses to ethical perspectives from Week 1 and Week 15 (the final week of the course) reveals that the course has influenced students towards more ethically aware stances concerning justice in education, critical consciousness (in the form of self-efficacy), and the intersection of technology with labor. These findings advocate for the continuation and expansion of ethically-focused curricula in computing disciplines, aiming to equip future professionals with the necessary tools to navigate and shape the technological landscapes with ethical integrity and social responsibility.

Our analysis of the survey data delved into the application of ethical concepts in relation to students' future. Survey aside from likert-scale questions asked about students' career intentions and industries of interest, linking their ethical perspectives to potential professional environments. In the data analysis phase, responses categorized under "others" were excluded due to their minimal frequency, thereby streamlining the analysis process. Additionally, questions formatted to elicit multiple responses were quantitatively analyzed based on their frequency counts. This approach ensured a more focused and efficient analysis of the predominant trends within the dataset. This multi-dimensional approach was designed to provide further understanding of how ethical discussions in an academic setting might influence students' views and decision-making processes as they transition into professional roles [15].

6.1 Ethical Framework

The stacked bar graphs in Figures 7 and 8 explore the shifts in student attitudes toward utilitarian outcomes and individual rights, respectively. Unintuitively, despite the inherent tension between a utilitarian and rights lens, students in Week 1 and Week 15 lean towards agreement for both statements. For the analysis, we will further examine *to what extent* the students support individual rights and liberties over utilitarian principles. From Figure 8, we observe an increased agreement with the statement "It is important to uphold the rights and liberty of a person, even if it means worse outcomes overall," compared to the diverging views on the necessity of making rules that prioritize the best outcomes for the majority at the expense of individual rights. Tangibly, the utilitarian ethics lens saw an increase of ~23% in agreement for a total of 53% (since students could not respond "Neutral"), while the rights ethics lens saw an increase of ~25% in agreement for a total of 64%.

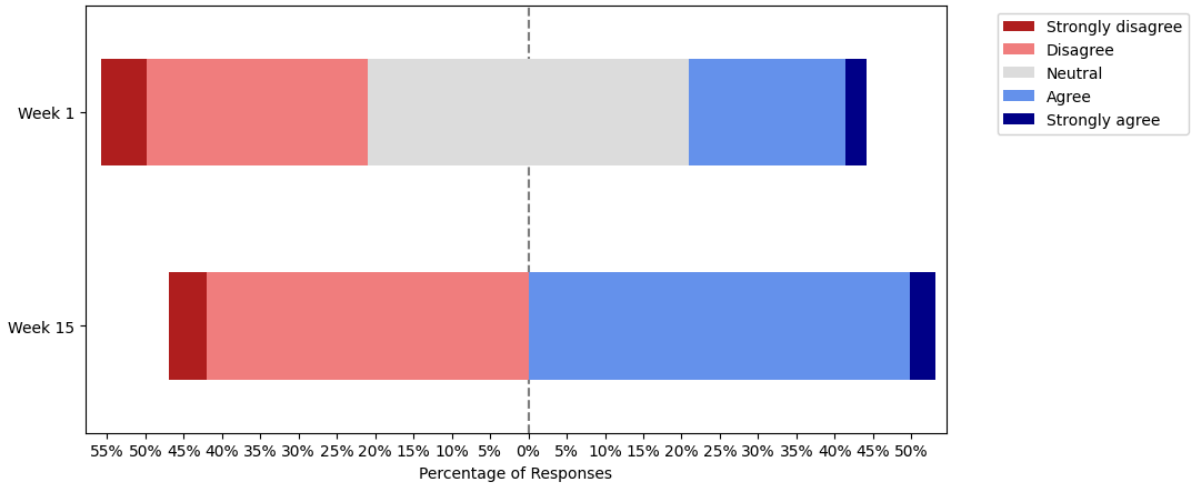

Figure 7: Diverging stacked bars to the survey question, "It is important to make rules based on what leads to the best outcomes for the most, even if it results in deprioritizing rights or liberties."

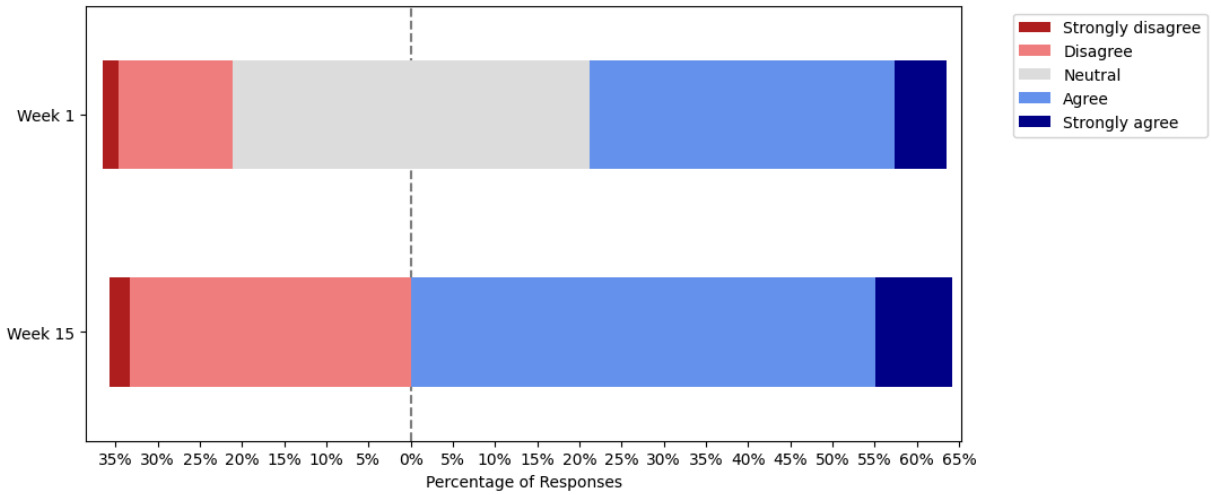

Figure 8: Diverging stacked bars to the survey question, "It is important to uphold the rights and liberty of a person, even if it means worse outcome overall."

6.2 Self-Efficacy

The analysis of student perceptions regarding the impact of regular citizens on enhancing their communities also exhibit shifts over the course of the semester. This question was primarily designed to analyze student self-efficacy. Here, we define self-efficacy as the perceived autonomy of a single individual to effect change on some system. This concept can then tie into an individual's critical consciousness fluency, as critical consciousness focuses particularly on power dynamics that repress populations. However, this question is limited by the fact that it

does not explicitly mention social systems or power. Notably, of the 35 students who initially remained neutral in Week 1, only 2 ended up responding with a negative response towards self-efficacy, whereas the remaining 33 responded with “Agree” or “Strongly agree” (Figures 9 and 10). This progression towards more affirmative stances indicates a positive shift in students’ attitudes towards community involvement and empowerment over the semester. The results reflect the course's emphasis on social responsibility and may also reflect broader pedagogical intentions to enhance critical consciousness through active civic participation and engagement.

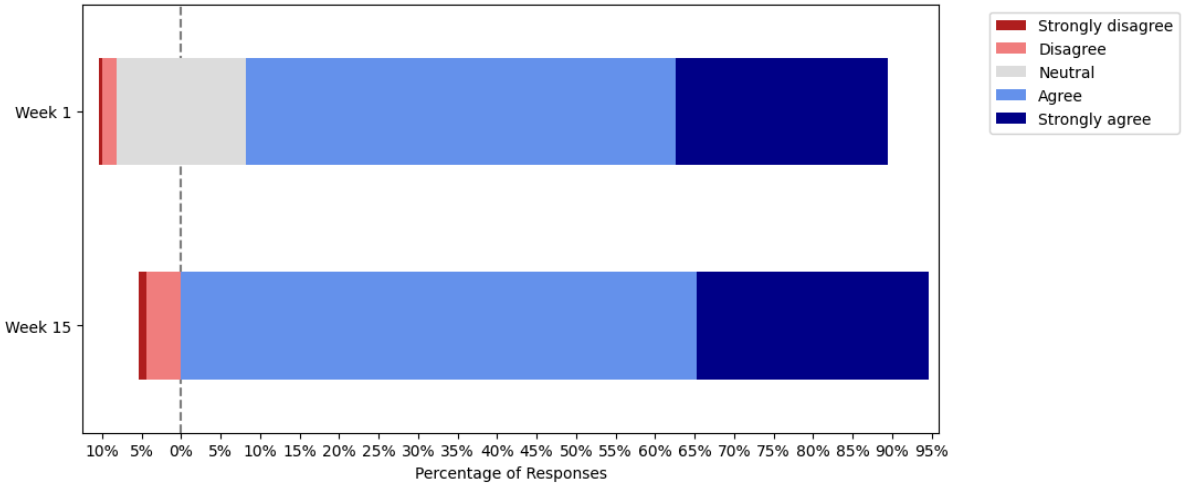

Figure 9: Diverging stacked bars to the survey question, “Regular citizens can have an impact in making my community a better place to live.

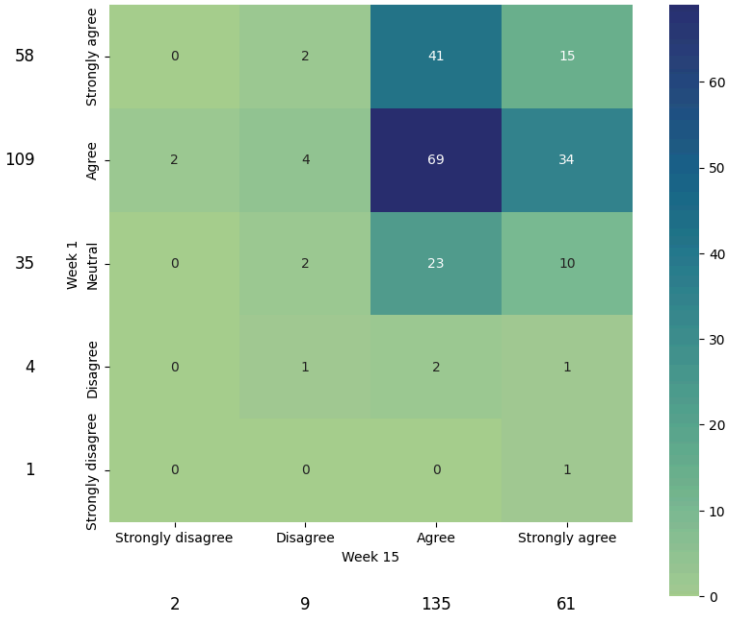

Figure 10: Cross tabulation to the survey question, “Regular citizens can have an impact in making my community a better place to live.

6.3 Tech and Labor

The analysis of the stacked bar graph in Figure 11 reveals a significant shift from neutrality to more definitive stances in favor of the recent local academic employee labor movement (negatively valenced responses were in favor of the Fall 2022 UC-wide academic strike). Of the 64 respondents who were initially neutral, 40 of them moved towards disagreement, while 24 of them agreed with the assessment that the labor strike achieved net harm. This is reflected as the diverging stacked bars shift to the left, given that the underlying assumption was that the neutral respondents are equally distributed on either side.

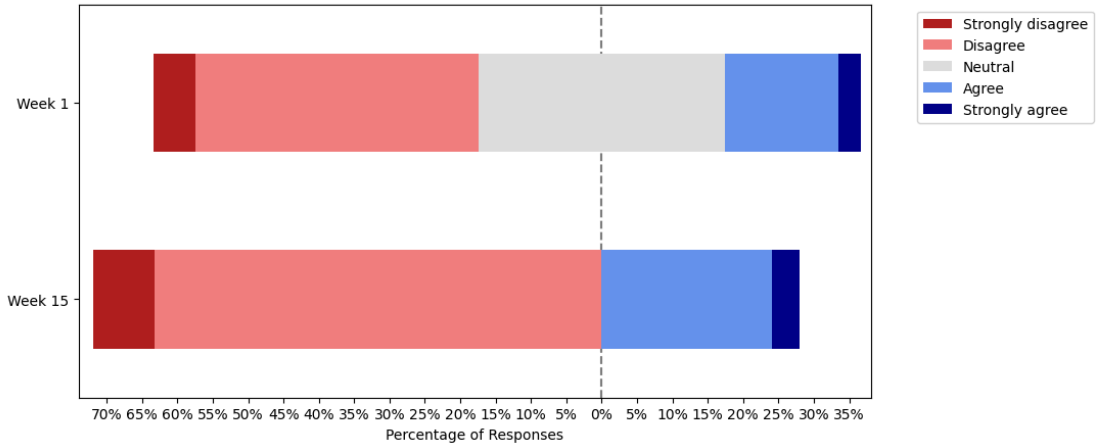

Figure 11: Diverging stacked bars to the survey question, "I think the UC-wide academic workers strike in Fall 2022 did more harm than good."

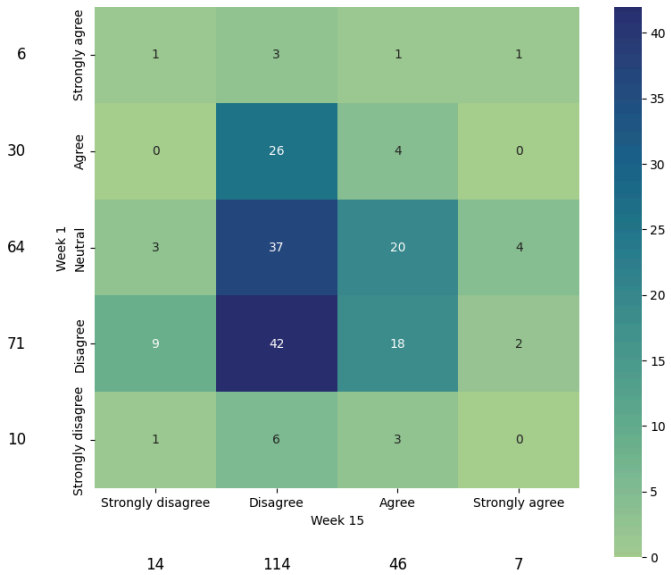

Figure 12: Cross tabulation to the survey question, "I think the UC-wide academic workers strike in Fall 2022 did more harm than good."

In reflecting on the data from diverging stacked bars and cross-tabulations, students' inclinations towards upholding individual rights, even at the expense of utilitarian outcomes, have strengthened. This shift could be attributed to the comprehensive coverage and case studies introduced through weekly readings, as well as encouragement of student engagement through QQCRs. The increase in agreement with the importance of upholding rights suggests that the course effectively highlighted the significance of protecting individual liberties within the tech landscape. Furthermore, the increase in students expressing a desire to engage in community building and civic responsibilities points to a successful installation of critical consciousness through the course.

These findings have demonstrated that ethics education in computing can influence students' ethical stances and potentially guide their future professional behaviors and decisions. The observed shifts towards a stronger endorsement of individual rights and increased critical consciousness reflect a promising integration of ethical reasoning into the students' professional ethos.

7 Future Work and Conclusion

For future work, we would like to continue iterating on the model we used to code the QQCR responses to ask more granular research questions in order to track specific QQCR responses. Additionally, we did not end up doing any concerted analysis controlling on graduation year or gender, which would provide more light on how this course can promote student belonging and inclusivity. The largest confounding factor in our survey analysis was that students could not remain neutral in the Week 15 survey—a choice intentionally made by the instructor for pedagogical purposes. As a result, respondents who chose to remain neutral after the term ended could not express that inclination given survey restrictions. Our research also centered around the usage of QQCRs in CS195, but we also collected QQCR data for CS H195 students. It would be interesting to compare the difference between QQCRs and survey data between CS195 and CSH195 students. Lastly, we'd like to investigate the ways that peer-to-peer interactions on the course forum impacts the sense of belonging, particularly in minoritized groups. As we set our sights on belonging, it's worth looking into other course mechanisms that can be adjusted to promote belonging, including the way that we approach discussion sections (Appendix IV).

Usage of Machine Learning Models for Classification. When we did our research on related work, we found that most published research surrounding computing ethics used either survey data (69% of computing ethics papers [39]) or a qualitative analysis of the utility of an assignment by examining student submissions [35]. Within this report, we attempted to scale the qualitative analysis by using modern fine-tuning techniques to generalize some work. However, much more work needs to be done in developing these models before developing more certainty on this method as a rigorous way of performing qualitative research using a codebook. We see this as a promising mechanism to evaluate students' submitted work at scale in the future.

Forum Management. Within our longitudinal analysis of a student's QQCRs, we encountered a student who had comments removed by staffers, which brings us to an important aspect of QQCRs: forum management. Since the nature of QQCRs necessitates interactions between students, staff had to consider a careful balance between valuing exposure (ensuring that students are able to see the full diversity of viewpoints that we leverage at Cal) while also holding students accountable to the community guidelines (Course Structure). Ultimately, we recognize that QQCRs are able to work at scale without increasing staff overhead linearly, but that does not mean that staff should completely neglect forum management.

Leveraging the scalability of discussion forums like EdStem to facilitate discussion among students enables them to critically engage with the material and gain exposure to the perspectives of their peers. By selecting topics and readings that contain a wide variety of culturally relevant perspectives, students are exposed to the concepts of cultural humility and critical consciousness via the stories and lessons shared. We observed a peak in student interaction (via the QQCR response length of their responses) with contemporary and immediately relevant topics like Generative AI (Weeks 7–8) and Tech and Labor (Weeks 11–15). Critical engagement was seen most frequently with topics that students were either already familiar with/found relatable, or expressed strong normative opinions on how the world should be (even if the students disagreed with the opinion expressed by the author). This finding does

not necessarily imply that the readings need to be densely packed with critical analysis (Section 5.2).

With that said, the lateness policy, flexibility in the specific readings that students could choose to engage with, and the number of deliverables all greatly impacted the length and distribution of student responses. Thus, we reaffirm the above claim: Although QQCRs enable students to engage critically with the material, there are a multitude of confounding factors that impact student engagement once disaggregated.

Finally, we found that student sentiment surrounding their ethical framework, critical consciousness, and tech and labor may have been influenced by the class. Given the cross tabular analysis and diverging stacked bars, we found that students generally became more justice-centered, realized their own self-efficacy to effect change, and recognized the systems around them that stifle change.

At this pivotal time of global, wide-scale change and crisis, we hope that this report has provided some promising first steps that address the needs of undergraduate computing curricula in higher education. The early frameworks and findings of this report can provide the classroom cultural spaces and interdisciplinary coursework to empower our computing students and promote belonging, equity, and social justice for all.

References

1. ACM Code of Ethics and Professional Conduct, "Preamble and Sections 1-4," Association for Computing Machinery, June 22, 2018. [Online]. Available: https://www.acm.org/code-of-ethics
2. Allison Bradford et. al. 2023. Adaptive dialog to support student understanding of climate change mechanisms and who is most impacted. In Blikstein, P., Van Aalst, J., Kizito, R., & Brennan, K. (Eds.), Proceedings of the 17th International Conference of the Learning Sciences - ICLS 2023 (pp. 816-823). International Society of the Learning Sciences.
3. Alvaro M. Bedoya. 2020. The Cruel New Era of Data-Driven Deportation. Slate. https://slate.com/technology/2020/09/palantir-ice-deportation-immigrant-surveillance-big-data.html [Accessed May 17, 2024].
4. Amy Orben and Sarah-Jayne Blakemore. 2023. How social media affects teen mental health: a missing link. Nature. https://www.nature.com/articles/d41586-023-00402-9 [Accessed May 17, 2024]
5. Anne G. Applin. 2006. A learner-centered approach to teaching ethics in computing. In Proceedings of the 37th SIGCSE technical symposium on Computer science education (SIGCSE '06). Association for Computing Machinery, New York, NY, USA, 530–534. https://doi.org/10.1145/1121341.1121505
6. Arno K. Kumagai MD; Monical L. Lypson MD. Beyond Cultural Competence: Critical Consciousness, Social Justice, and Multicultural Education. Academic Medicine 84(6):p 782-787, June 2009. | DOI: 10.1097/ACM.0b013e3181a42398
7. Barbara J. Grosz, David Gray Grant, Kate Vredenburgh, Jeff Behrends, Lily Hu, Alison Simmons, and Jim Waldo. 2019. Embedded EthiCS: integrating ethics across CS education. Commun. ACM 62, 8 (July 2019), 54--61. DOI: https://doi.org/10.1145/3330794
8. Benjamin Bloom. 1984. The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring. *Educational Researcher,* Vol. 13, No. 6 pp. 4-16.
9. Beleicia B. Bullock, Fernando Nascimento, and Stacy A. Doore. 2021. Computing Ethics Narratives: Teaching Computing Ethics and the Power of Predictive Algorithms. In Proceedings of 52nd Technical Symposium on Computer Science Education (SIGCSE'21), March 13-20, Virtual Event USA, ACM, NY, NY, USA. 7 pages. https://doi.org/10.1145/3408877.3432468
10. Ben Tarnoff. 2020. The Making of the Tech Worker Movement. Logic(s). https://logicmag.io/the-making-of-the-tech-worker-movement/full-text/ [Accessed May 17, 2024]
11. Bill Cowan, Elodie Fourquet, Marta Kryven. 2018. Teaching the Societal Consequences of Computer Science: New Ideas for Increasing Student Involvement. In Proceedings of 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE'18). ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3197091.3197146
12. California Assembly Bill 2097. 2024. https://legiscan.com/CA/bill/AB2097/2023 [Accessed May 17, 2024]
13. Carolyn Chen. Work Pray Code: When Work Becomes Religion in Silicon Valley. Princeton University Press, 2022. Project MUSE. muse.jhu.edu/book/99909.

14. Chima Abimbola Eden, Onyebuchi Nneamaka Chisom, Idowu Sulaimon Adeniyi. 2024. Cultural Competence in Education: Strategies for Fostering Inclusivity and Diversity Awareness. International Journal of Applied Research in Social Science. 10 pages. https://doi.org/10.51594/ijarss.v6i3.895
15. Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In The 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20), March 11–14, 2020, Portland, OR, USA. ACM, New York, NY, USA, 7 pages. https://doi.org/10. 1145/3328778.3366825
16. CS 195, Social Implications of Computing Course Website. 2005. http://wla.berkeley.edu/~cs195/sp05/ [Accessed May 17, 2024]
17. CS 195, Social Implications of Computing Course Website. 2023. https://fa23.cs195.org/assignments/podcast1.html [Accessed May 17, 2024]
18. CS61B, Data Structures Course Website. 2021. https://inst.eecs.berkeley.edu/~cs61b/fa21/ [Accessed May 17, 2024]
19. Continuum of Cultural Competency. 1989. https://servicelearning.ctb.ku.edu/wp-content/uploads/ContinuumOfCulturalCompetency-1024x576.jpg [Accessed May 17, 2024]
20. D.J. Carter. (2008). Cultivating a critical race consciousness for African-American school success. *Educational Foundations, 22* (1-2), 11-28.
21. EECS Major Upper Division Degree Requirements. 2024. https://eecs.berkeley.edu/resources/undergrads/eecs-2/degree-reqs-upperdiv-2/ [Accessed May 17, 2024]
22. Embedding Ethics in Computer Science. 2024. https://embeddedethics.stanford.edu/ [Accessed May 17, 2024]
23. Emanuelle Burton, Kristel Clayville, Stacy A. Doore, Michael S. Kirkpatrick, and Michael Goldweber. 2024. Managing Authority When Teaching Computing Ethics. In Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2 (SIGCSE 2024). Association for Computing Machinery, New York, NY, USA, 1523–1524. https://doi.org/10.1145/3626253.3631659
24. Emily Vogels and Risa Gelles-Watnick. 2023. Teens and social media: Key findings from Pew Research Center surveys. Pew Research Center. https://www.pewresearch.org/short-reads/2023/04/24/teens-and-social-media-key-findings-from-pew-research-center-surveys/ [Accessed May 17, 2024]
25. Francesca Paris, Larry Buchanan. 2023. 35 Ways Real People Are Using A.I. Right Now. The New York Times. https://www.nytimes.com/interactive/2023/04/14/upshot/up-ai-uses.html [Accessed May 17, 2024]
26. Gianluca Russo. 2022. For Plus-Size Creators, Tiktok Presents a New Wave of Challenges. Nylon. https://www.nylon.com/life/tiktok-body-positivity-plus-size-creators [Accessed May 17, 2024]
27. Grace Barkhuff, Jason Borenstein, Daniel Schiff, Judith Uchidiuno, and Ellen Zegura. 2024. Considerations for Improving Comprehensive Undergraduate Computing Ethics Education. In Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2 (SIGCSE 2024), March 20– 23, 2024, Portland, OR, USA. ACM, New York, NY, USA, 2 pages. https: //doi.org/10.1145/3626253.3635557

28. Helen Noble, Gary Mitchell. What is grounded theory? *Evidence-Based Nursing* 2016;**19:**34-35.
29. Huang, X. (2017). Example-based learning: Effects of different types of examples on student performance, cognitive load and self-efficacy in a statistical learning task. Interactive Learning Environments, 25, 283–294. http://dx.doi.org/10.1080/ 10494820.2015.1121154
30. Jim Ormond. 2024. With Great Computing Power Must Come Responsible Computing. ACM Journal on Responsible Computing. https://www.acm.org/media-center/2024/april/inaugural-jrc [Accessed May 17, 2024]
31. Jodi Kantor and Arya Sundaram. 2022. The Rise of the Worker Productivity Score. The New York Times. https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html [Accessed May 17, 2024]
32. Kalir, J., Morales, E., Fleerackers, A., & Alperin, J. (2020). "When I saw my peers annotating:" Student perceptions of social annotation for learning in multiple courses. Information and Learning Sciences, 121(3/4), 208; see also Novak, E., Razzouk, R., & Johnson, T. E. (2012). The educational use of social annotation tools in higher education: A literature review. The Internet and Higher Education, 15(1), 39-49.
33. Kate Vredenburgh, Jeff Behrends, Lily Hu, Alison Simmons, and Jim Waldo. 2019. Embedded EthiCS: integrating ethics across CS education. Commun. ACM 62, 8 (July 2019), 54--61. DOI: https://doi.org/10.1145/3330794
34. Kevin Lin. 2022. CS Education for the Socially-Just Worlds We Need: The Case for Justice-Centered Approaches to CS in Higher Education. In Proceedings of the 53rd ACM Technical Symposium on Computer Science Education - Volume 1 (SIGCSE 2022), Vol. 1. Association for Computing Machinery, New York, NY, USA, 265–271. https://doi.org/10.1145/3478431.3499291
35. Mara Kirdani-Ryan and Amy J. Ko. 2022. The House of Computing: Integrating Counternarratives into Computer Systems Education. In Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 1
36. Matthew Diemer, Luke Rapa, Adam Voight, Ellen Mcwhirter. (2016). Critical Consciousness: A Developmental Approach to Addressing Marginalization and Oppression. Child Development Perspectives. 10. 10.1111/cdep.12193.
37. Naomi Klein. 2023. AI machines aren't 'hallucinating'. But their makers are. The Guardian. https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomi-klein [Accessed May 17, 2024]
38. Nick Anderson. 2023. College is remade as tech majors surge and humanities dwindle. Washington Post. https://www.washingtonpost.com/education/2023/05/19/college-majors-computer-science-humanities/ [Accessed May 17, 2024]
39. Noelle Brown, Benjamin Xie, Ella Sarder, Casey Fiesler, and Eliane S. Wiese. 2024. Teaching Ethics in Computing: A Systematic Literature Review of ACM Computer Science Education Publications. ACM Trans. Comput. Educ. 24, 1, Article 6 (January 2024), 36 pages. https://doi.org/10.1145/3634685
40. Paul Freire. (1970). *Pedagogy of the oppressed.* New York, NY: Herder & Herder.

41. Pamela Samuelson, Generative AI meets copyright. *Science* **381**,158-161(2023). DOI:10.1126/science.adi0656
42. Rob Reich, Mehran Sahami, Jeremy M. Weinstein, and Hilary Cohen. 2020. Teaching Computer Ethics: A Deeply Multidisciplinary Approach. In Proceedings of the 51st ACM Technical Symposium on Computer Science Education (Portland, OR, USA) (SIGCSE '20). Association for Computing Machinery, New York, NY, USA, 296–302. https://doi.org/10.1145/3328778.3366951
43. Ruha Benjamin. *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity, 2019.
44. Safiya Noble. 2018. Algorithms of oppression: How search engines reinforce racism. New York University Press.
45. Sharon Goldman. 2023. Open letter calling for AI 'pause' shines light on fierce debate around risks vs. hype. VentureBeat. https://venturebeat.com/ai/open-letter-calling-for-ai-pause-shines-light-on-fierce-debate-around-risks-vs-hype/ [Accessed May 17, 2024]
46. Y.-T. C. Yang (2008). A catalyst for teaching critical thinking in a large university class in Taiwan: Asynchronous online discussions with the facilitation of teaching assistants. Educational Technology Research and Development, 56, 241–264. http://dx.doi.org/10.1007/s11423-007-9054-5
47. Z. Szabo, J. Schwartz. (2011). Learning methods for teacher education: The use of online discussions to improve critical thinking. Technology, Pedagogy and Education, 20, 79–94. http://dx.doi .org/10.1080/1475939X.2010.534866
48. A Framework for Ethical Decision Making. 2021. Markkula Center for Applied Ethics at Santa Clara University. https://www.scu.edu/ethics/ethics-resources/a-framework-for-ethical-decision-making/ [Accessed May 17, 2024]

Appendix I: CS195 Fall 2023 Syllabus

Course Descriptions and Enrollment

CS 195 is a discussion-intensive course about the social implications of computer technology. The purpose of this course is to help students make informed and thoughtful choices about their careers, participation in society, and future development activities. Readings and lecture topics are drawn from a range of fields that together seek to describe our contemporary global society: sociology, philosophy, economics, public policy, etc.

- Course Catalog
- If you are looking to enroll in this course but do not meet EECS department enrollment categories, please make a private post on Ed.

CS H195 discussions complement and supplement the topics presented in CS 195 lecture series. Students are expected to engage at a deeper level with the assigned weekly readings and be prepared to engage in thoughtful and constructive discussions around the course material. Where applicable, we will invite guest speakers from industry and academia to present their work and to participate in discussions alongside students. CS H195 will encourage students to go beyond thinking about computer technology as solely an engineering problem but instead viewing it holistically from the perspective of social sciences, legal studies, policymaking, equity, and inclusion.

CS H195 is a 3-unit class with the following expectations:

- In addition to CS 195 lecture, you attend 90-minute in-person discussion sections led by a GSI, once weekly.
- You will actively discuss each CS 195 topic in more depth.
- You will work on interesting projects / papers that replace the usual CS 195 assignments.
- If interested, please fill out the interest form listed on the Course Catalog.
- **Course applications for H195 are due Friday 9/1 11:59pm PT.**
- The discussion page has the discussion schedule and syllabus.

Staff Contact

The best (and fastest) way to contact the staff is through the Ed discussion board. Additionally, feel free to drop by or schedule a meeting for our office hours. More information about the office hours is available in the left navigation panel.

If you need to contact the course staff for personal reasons, we can be reached at email addresses in the sidebar. All emails end with **berkeley.edu**

Course Components and Grading

CS 195 is graded P/NP. The reason for this policy is to ensure that you can feel free to express opinions that are different from those of the instructors, both in class meetings and in written work.

In order to receive a passing grade in CS 195, you must complete the following:

Course Component	Overview
Update 10/11 Lectures	Attend at least 9 out of 12 lectures (excludes Lecture 01 and 08).
Surveys	Complete all course surveys (~3 in the semester).
Readings / QQCR	Submit a QQCR on each set of 13 lecture readings (excludes Lecture 01).
Update 9/28 Podcasts	Submit two Computing in the News podcasts that receive passing peer review grades. Also, provide 2 peer reviews for each podcast.

Lectures

Lecture will be each Wednesday 5:00-6:30PM hybrid in Evans 10 with a Zoom simulcast. Participation will be expected for both in-person and Zoom students.

Lecture is going to be a mix of presentations, students asking questions, small group discussion, and class-wide debriefing. Participation will be expected.

Attendance will be taken in class and is required either online or in-person. You cannot have a time conflict. The lecture is the vast majority of the experience of the class. There will *not* be after-the-fact recordings posted.

Surveys

There will be periodic surveys throughout the semester, announced during class and on Ed. Please submit the survey on time by the date listed on the homepage.

- Pre-semester survey: Lecture 01
- Technology and identity survey: Lecture 06
- End-of-semester survey: Lecture 15 (note: this is separate from campus course evaluations)

Readings / QQCR

Readings for each week will be released the Friday before each next class (or thereabouts). **Please finish the readings before class!**

QQCR, or Question, Quote, Comment, and Reply, is a short weekly assignment on the readings to help you reflect on the readings *and* contribute to an ongoing class discussion. (Updated 8/30 to reduce workload.)

To complete a QQCR,

- Post a question, quote, comment, or reply on the associated Ed megapost for a reading.

    - **Question**: Make a new thread under the megapost with a question about the reading and any needed context for others to follow up. You should include the text [Question] at the beginning of your post.
    - **Quote**: Make a new thread under the megapost with a quote from the reading, and your comment/reaction to that quote. You should include the text [Quote] at the beginning of your post.
    - **Comment**: Make a new thread under the megapost with a comment about the reading. You should include the text [Comment] at the beginning of your post.
    - **Reply**: Make a response on *someone else*'s thread in the associated Ed megapost for a reading.
- It is OK for you to post anonymously, meaning your name will be anonymous to fellow classmates but visible to course staff (who will use it for tracking assignment credit).
- See below for grading and community guidelines.

To receive full credit for each set of lecture readings, you must do two things:

- (1) You must make four **two** of a Question, Quote, Comment, Reply on the associated Ed megapost(s) for the required or recommended readings.
- (2) You must **also** make a Reply on someone else's thread on one of the associated Ed megapost(s) for one of the required or recommended readings
- Restated, you must make five **three** QQCRs *total* on *any* of the required/recommended readings, where one of your QQCRs needs to be a reply.
- (1) and (2) not need to be on the same reading. However, they should be made on readings listed as "Required" or "Recommended" for that week. QQCRs made for "Extra" or "Optional" readings do not count for course credit. For example:
    - A QQCR could consist of a mix of three two Questions on one required reading, and two Replies one Reply on a different recommended reading.
    - Or, a QQCR could consist of five three Replies on the same recommended reading. These Replies can be in the same thread.
- See this Ed post for clarifications.

QQCR Grading:

- Again, the main goal of this assignment is to give you a chance to reflect on readings and contribute to discussion. We expect there to be a range of opinions in this classroom. As a result, **one-word posts will not receive credit.** Most other posts will.
- You *should* post in the **Ed megapost** that course staff create for each reading.
- You *should not* make private or separate public Ed posts. Doing so will not earn credit. Early on, we might make an effort to merge your posts to the respective Ed megapost, but ultimately you are responsible for moving your posts to the megapost.
- "Heart"s on other student posts are encouraged but unfortunately do not count for credit (we have no way of tracking hearts).
- **Deadline**: You should complete the QQCR for a lecture prior to the lecture itself. See the Accommodations and Late Policy section below for more details.

We'll do our best to release grades for the first set of QQCRs early, so that you can understand the grading scheme and update and/or correct your QQCR submissions.

Community guidelines

Our classroom environment extends to online spaces like Ed. Be civil and respect others, even those that are not in this course. There are 3 ways for you to address hateful comments and

prevent harassment, in order of decreasing emotional labor (the last two do not count towards QQCR):

- Reply on Ed directly to post.
- Submit an anonymous report form to course staff: https://forms.gle/bBmqFwRpk3wxdoXn7.
- Reply on Ed with "That's not cool."

Course staff reserve the right to delete and/or not award credit or dismiss students from seminar for remarks that promote hate speech, trolling, harrassment, discrimination, or defamation. Serious offenders will be reported directly to both the CS Department Grievances faculty member (link) and the University Center for Student Conduct.

Podcasts

Update 9/28 In lieu of the typical essays for this course, this semester you will submit two podcasts of varying format in relation to issues discussed in class. Please see the home page for more details.

Podcasts are peer reviewed, which means you will review **two podcasts** per peer review window, and have each of your podcasts reviewed by two of your peers.

There will be no essays this semester; furthermore we expect the podcasts to actually take up less time than formulating your thoughts on paper. Please do contact us if this is not the case.

Accommodations and Late Policy

We understand that situations come up during the semester. While we will make an effort to contact students about missing work throughout the semester, *it is ultimately your responsibility to follow up with us to request additional accommodations as needed*.

Lecture attendance: Please use your built-in absences first before contacting staff about additional makeups. You may make up lecture attendance by attending the H195 discussion section (discussion meeting times in sidebar). Note that sections often require prework, so please do make a private post on Ed in advance so we can provide you with details.

Surveys: You are expected to complete all surveys by their listed deadlines on the homepage. These will be reopened for completion as needed. To request a reopening, please make a private post on Ed.

Readings/QQCR: We prefer that you complete each QQCR by the *listed lecture date* to ensure a rich discussion during class. However, if you cannot make this date, you will have until the **last day of classes** to submit all outstanding QQCRs. Please submit regrades via the form below.

Podcasts: *As of now there are no plans to offer makeup podcasts.* While we do not want this 1-unit seminar to stress you out, please also understand that because of our grading timelines, podcasts must be submitted by the posted deadline. If you have extenuating circumstances or anticipate an overlap with your commitments (personal emergencies, exams, etc.), please make a private Ed post as soon as possible or see the form below so we can provide a short extension or another accommodation. We will review each extension request on an individual basis.

DSP: If you have a DSP-approved accommodation, we will contact you by the end of Week 2 with this semester's DSP policies. Requesting assignment extensions according to your letter will involve you making a private post on Ed *prior* to the assignment deadline.

If you need an extension or regrade request, please fill out this semester's Extension/Regrade Request form.

Academic Integrity

The student community at UC Berkeley has adopted the following Honor Code: "As a member of the UC Berkeley community, I act with honesty, integrity, and respect for others." The expectation is that you will adhere to this code.

Plagiarism/Self-plagiarism: You must be original in composing the writing assignments in this class. To copy text or ideas from another source (including your own previously, or concurrently, submitted course work) without appropriate reference is plagiarism and will result in a failing grade for your assignment and usually further disciplinary action. Anyone found to have submitted assignments completed by another person (student or non-student) will likewise receive a failing grade in the course and be reported to the Center for Student Conduct. For additional information on plagiarism, self-plagiarism, and how to avoid it, see this GSI Guide for Preventing Plagiarism.

Appendix II: Readings and Codes

Week 1

No required readings

Week 2

1. Pew Research Center, "Teens and social media: Key findings from Pew Research Center surveys". 2023.
    a. SS-SV
2. Nature. "How social media affects teen mental health: a missing link". 2023.
    a. J-S, CR-L, SS-SV
3. NPR. "New Research Casts Doubt On Connection Between Smartphone Use And Teen Mental Health". 2019.
    a. J-S, CR-L, SS-SV
4. Nylon. "For plus-size creators, Tiktok presents a new wave of challenges". 2022.
    a. J-S, CR-P, SS-SV, ER-J, ER-R, C, P-I, CC-R

Week 3 - Data Privacy

1. The New York Times, "The Rise of the Worker Productivity Score", 2022.
    a. J-S, ER-R, CR-P
2. EU, "What is GDPR, the EU's new data protection law?", 2022. Wired. (references child abuse).
    a. I, J-S, ER-R
3. The New York Times, "Are You in a Gang Database?", 2020.
    a. ER-R, ER-J, J-S, CR-P, P-A, CC-R, SS-POL

Week 4 - Software Risks and Professional Ethics

1. Fabio, "Killed by a Machine: The Therac-25", 2015.
    a. J-S, ER-J, CR-P, CR-L, SS-SV
2. The ACM, "Code of Ethics", 2018.
    a. ER-R, ER-J, ER-CV, J-S, A-O, CR-E, SS-ED, I
3. Miller, "Designing Ethical Self-Driving Cars", 2023.
    a. ER-CV, J-S, A-O, CR-H, SS-SV

Week 5 - Technosolutionism

1. Benjamin, *Race After Technology*: Introduction, 2019.
    a. ER-R, ER-J, J-IR, A-O, CR-P, CR-L, C, P-A, CC-A, SS-FAM, SS-SV, SS-ED, SS-EC, SS-POL

Week 6 - Government Censorship and Surveillance

1. NPR, Surveillance And Local Police: How Technology Is Evolving Faster Than Regulation, 2021.
   a. ER-R, ER-J, ER-CV, J-IR, A-O, CR-P, CR-L, CC-R, SS-POL
2. Slate, “The Cruel New Era of Data-Driven Deportation”, 2020.
   a. ER-R, ER-J, ER-CV, J-S, CR-P, CR-L, CR-E, P-I, CC-R, SS-SV, SS-POL
3. NPR, “A decade on, Edward Snowden remains in Russia, though U.S. laws have changed”, 2023.
   a. ER-R, SS-POL

Week 7 - The AI Debate

1. VentureBeat. “Open letter calling for AI ‘pause’ shines light on fierce debate around risks vs. hype”, 2023.
   a. ER-CV, ER-U, J-IR, A-O, CR-P, CR-E, SS-SV, SS-POL, I
2. VentureBeat. “Titans of AI Andrew Ng and Yann LeCun oppose call for pause on powerful AI systems”, 2023.
   a. ER-U, J-IR, A-O, CR-L, CR-E, CR-H, SS-POL, SS-SV
3. Berkeley News. “UC Berkeley historian of science ponders AI’s past, present and future”, 2023.
   a. ER-CV, J-IR, J-S, A-O, CR-L, CR-E, P-A, SS-SV, SS-ED, SS-EC, SS-POL
4. Office of the Governor. “Governor Newsom Signs Executive Order to Prepare California for the Progress of Artificial Intelligence”, 2023.
   a. ER-CV, SS-ED, SS-POL

Week 8 - Google AI Ethics

1. Wired, “What Really Happened When Google Ousted Timnit Gebru”, 2021.
   a. Not coded (paywalled after the semester ended)
2. Platformer. “The withering email that got an ethical AI researcher fired at Google”, 2020.
   a. ER-J, J-S, A-O, CR-L, CR-P, P-A, SS-SV
3. Google Walkout For Real Change. “Standing with Dr. Timnit Gebru — #ISupportTimnit #BelieveBlackWomen”, 2020.
   a. ER-J, J-S, A-O, CR-P, CR-L, CR-E, P-I, SS-SV

Week 9 - Generative AI, Part II

1. Science, “Generative AI meets copyright”, 2023.
   a. ER-R, ER-U, J-S, CR-L, SS-SV, SS-POL
2. The NYT, “35 Ways Real People Are Using AI Right Now”, 2023.
   a. SS-FAM, SS-ED
3. Guardian. “AI machines aren’t ‘hallucinating’. But their makers are”, 2023.
   a. ER-R, ER-J, J-S, CR-P, CR-E, CR-L, SS-SV, SS-EC, SS-POL

Week 10 - Silicon Valley and the U.S.

1. NPR, “FTC Chair Lina Khan’s lawsuit isn’t about breaking up Amazon, for now”, 2023.
   a. ER-J, J-S, CR-P, CR-E, CR-L, SS-SV, SS-EC, SS-POL

2. KQED, "Court Upholds Prop. 22 in Big Win for Gig Firms Like Lyft and Uber", 2023.
    a. ER-R, J-S, SS-ED, SS-POL
3. Giblin and Doctorow, *Chokepoint Capitalism*: Chapter 1, 2022.
    a. ER-CV, ER-R, J-S, A-O, CR-P, CR-L, CR-E, CC-M, SS-SV, SS-EC, SS-POL

Week 11 - Tech and Labor

- Chen, *Work Pray Code*. Introduction: How Work Is Replacing Religion and Conclusion: Techtopia 2022.
    - ER-R, J-S, CR-P, CR-H, CR-L, SS-SV

Week 12 - CS Education

1. Sepehr Vakil, Jennifer Higgs. "It's about power", pg. 31-33, 2019.
    a. ER-J, J-IR, A-O, CR-L, CR-H, CC-A, SS-ED
2. Case Study 1
    a. Bradford, et. al. "Adaptive Dialog to Support Student Understanding of Climate Change Mechanism and Who is Most Impacted", 2023.
        i. J-IR, CR-L, SS-ED
3. Case Study 2
    a. Sal Khan, "Harnessing GPT-4, so that all students benefit. A nonprofit approach for equal access", 2023.
        i. ER-J, J-IR, A-O, CR-P, SS-ED
    b. Open AI, "Customer Stories: Khan Academy", 2023.
        i. ER-J, J-IR, A-O, CR-P, SS-ED

Week 13 - Memes, Misinformation, and Media Literacy

1. Phillips and Milner, *You Are Here*: Chapter 2, The Root of All Memes, 2021.
    a. ER-CV, J-IR, A-O, CR-H, CR-L, CR-P, P-V, CC-R, SS-FAM
2. Tanksley, "When Black Death Goes Viral: How Algorithms of Oppression (Re)Produce Racism and Racial Trauma", 2023.
    a. ER-R, ER-J, J-S, CR-P, CR-L, CC-R, SS-EC, SS-SV

Week 15 - Tech and Labor Part II

1. Logic Magazine. "The Making of the Tech Worker Movement.", 2020.
    a. ER-R, ER-J, ER-CV, J-S, A-O, CR-E, CR-H, CR-L, P-I, CC-A, SS-SV, SS-EC, SS-POL
2. TechCrunch. "East Of Palo Alto's Eden: Race And The Formation Of Silicon Valley", 2015.
    a. ER-R, ER-J, ER-CV, J-S, A-O, CR-P, CR-L, C, CC-R, SS-FAM, SS-SV, SS-ED, SS-EC, SS-POL

Appendix III: Full List of Average Number of Codes Applied to QQCRs Grouped by Reading

Title	Average Number of Codes Applied to QQCR Response	Number of Codes Applied to Reading
Case Study 1: Use of AI in education in a research context	2.671533	3
Nature: How social media affects teen mental health	2.746667	3
Pew Research: Teens and social media	2.782738	1
Logic Magazine: The Making of the Tech Worker Movement	2.804487	13
Carolyn Chen: Work Pray Code, Introduction	2.834286	6
Case Study 2: Use of AI in education in a commercial/ed tech context	2.837321	5
Carolyn Chen: Work Pray Code, Conclusion	2.867159	6
KQED: Court Upholds Prop. 22 in Win for Lyft, Uber	2.899543	4
Sepehr Vakil, Jennifer Higgs: "It's about power" pg. 31-33	2.912727	7
Platformer: The email that got an AI researcher fired	2.916279	7
NPR: Doubts On Connection Between Phone Use And Teen Mental Health	2.928854	3
NPR: FTC Chair Lina Khan brings lawsuit against Amazon	2.952218	8
Google Walkout For Real Change: Standing with Dr. Gebru	2.965909	8

Ruha Benjamin: Race After Technology	2.976628	14
VentureBeat: Open letter calling for AI ‘pause’	3.009259	8
TechCrunch: East of Palo Alto's Eden	3.022013	14
VentureBeat: Titans of AI Andrew Ng and Yann LeCun oppose call for pause	3.044304	8
ACM: Code of Ethics	3.057554	7
Stanford: Designing Ethical Self-Driving Cars	3.062937	5
Hackaday: Killed by a Machine, The Therac-25	3.090090	5
EU: What is GDPR?	3.097345	3
Giblin and Doctorow: Chokepoint Capitalism, Chp 1	3.098361	11
The Guardian: AI machines aren't hallucinating, their makers are	3.099057	9
Berkeley News: Historian of science ponders AI's past, present and future	3.105263	11
Tanksley: When Black Death Goes Viral	3.130102	8
The NYT: Are You in a Gang Database?	3.133333	7
NPR: A decade on, Edward Snowden remains in Russia, though U.S. laws have changed	3.147619	2
Office of the Governor: Executive Order to Prepare California for the Progress of Artificial Intelligence	3.148936	3

Wired: What Really Happened With Timnit Gebru	3.176230	N/A, paywalled
Phillips and Milner, You Are Here: Chapter 2	3.225000	9
NPR: Surveillance And Local Police: How Technology Is Evolving Faster Than Regulation	3.251163	9
Slate: The Cruel New Era of Data-Driven Deportation	3.321782	11
The NYT: 35 Ways Real People Are Using AI Right Now	3.383621	2
The NYT: The Rise of the Worker Productivity Score	3.493976	3
Science: Generative AI meets copyright	3.497382	6
Nylon: For plus-size creators, Tiktok presents a new wave of challenges	4.812903	8

Appendix IV: Bridge Discussions

Considering the scale that we operate at in Berkeley, educators here have to constantly think about scale. Unfortunately, policies that operate at scale oftentimes fail to seriously consider the ramifications on students whom that policy cannot properly cater to. We see this problem particularly with regards to feelings of belonging, thus feeding into the tropes of rampant imposter syndrome and isolation at Berkeley. Simultaneously, with so many students, it can be difficult to cater instruction to meet each student's needs in various aspects of the course. Thus, we promote bridge discussions as a potential mechanism to address both concerns.

Discussion sections are typically considered the hub for active learning — once a student has had their initial exposure to the material within the context of lecture, they can solidify their conceptual understanding with a teaching assistant (TA) via a once a week, 50-minute section. Here's the issue: in a class with 2,000 students, a given discussion section could have wildly disparate skill levels before even entering the course. Paired with the time constraints, TAs often find themselves having to move on well before all the students are able to attain a sufficient understanding of the material. In addition, traditional discussion sections can have up to a 30:1 student-to-TA ratio, and receiving specialized support can be difficult. One final issue is that this conventional model makes a crucial assumption: that students have been able to watch/attend the relevant lecture(s) before coming to discussion. We won't attempt to refute the pedagogical value of the traditional expectations of students, but rather, we introduce bridge discussions as a way to support our students who may not have the bandwidth to participate in our courses as they currently stand.

Thus, we draw inspiration from the format of smaller, liberal arts colleges (SLACs) to introduce bridge discussions as an alternative method of encouraging active learning within the class. We achieve this by combining the initial exposure of the material during lecture and the active learning aspect of discussion into one synchronous activity. Notably, these discussion sections cover the same exact material, but with double the amount of time and staff support. Bridge discussion attendees are also able to form deeper connections with their peers and TAs due to the structure of the section. By proactively targeting outreach for our bridge sections to students that may benefit from it, we are able to manage demand (even at such a scale) and equitably support all students without compromising on educational quality. Future work in this area should examine how the unique combination of an extended time discussion and having an increased staff-to-student ratio impacts student outcomes and sense of belonging.