2023 Annual Conference & Exposition
Baltimore Convention Center, MD | June 25 - 28, 2023
The Harbor of Engineering Education for 130 Years
ASEE
Paper ID #39306

Characterizing student argument justifications in small group sociotechnical discussions

Dr. Chelsea Joy Andrews, Tufts University

Chelsea Andrews is a Research Assistant Professor at Tufts University, at the Center for Engineering Education and Outreach (CEEO).

Ms. Fatima Rahman, Tufts Center for Engineering Education and Outreach

STEM Education graduate student at Tufts University

Characterizing student arguments against a technology in small group sociotechnical discussions

Introduction

Universities across the U.S. are finally heading the many calls to include sociotechnical thinking–grappling with issues of power, history, and culture–throughout the undergraduate engineering curriculum. While non-purely-technical topics have historically been relegated to separate courses, universities are now working to integrate sociotechnical content in courses previously considered to be purely technical. Researchers have varying motivations for this focus, including to better prepare students for engineering practice, which is inherently sociotechnical [1]; to increase the sense of belonging of historically excluded students, who are more likely to be interested in the social aspects [2]; and to create better societal outcomes [3-5]. Attempts to disrupt the social/technical dualism have included revising stand-alone ethics courses and adding more social components to previously purely technical courses, such as design courses [6-9].

Research in this space is still identifying what to expect of students and how to support deeper engagement in sociotechnical topics. This is being investigated through, for example, analyzing student interviews and focus groups [10-13], in class whole-group discussion [7, 14], and students' written work [15]. Here, we build on this research base by looking at small group in-class discussions.

This study is part of an NSF-funded research project to implement and study integrating sociotechnical components throughout a first-year computing for engineers course. In one iteration of the revised course, each week students read a news article on a current example of the uneven impacts of technology, then engaged in in-class small-group discussions supported by upper-class learning assistants. One of our project goals is to better understand these sociotechnical small group discussions: how to frame the topics, how to support students, and how to decide what makes discussions productive–including how to define productive in this case. In this first study on these discussions, we analyze two groups of students who uniformly rejected the technology under consideration–they did not uphold the common narratives related to technology. However, while both groups were critical of the technology, only one group grounded their arguments in concerns about harm. In this paper, we analyze students' discourse to answer the research questions: In small group discussions where students unanimously oppose the use of a technology, what arguments do they use to argue against that technology? How do these arguments relate to common narratives about technology?

Framework: Common narratives about technology

Prior research has worked to identify and describe the ways technology is commonly viewed in modern society. These uncomplicated, positive views are found in media [16], education standards [17, 18], and throughout the undergraduate engineering curriculum [19, 20]. We will use 'common narratives' to describe these perspectives, while others have used engineering mindsets, underlying worldviews [20], and ideologies [3]; these perspectives are also aligned with Radoff et al.'s description of narrow thinking [13]. Previous research has looked at how undergraduate engineering students align with these common narratives in interviews and focus groups with students [10, 13], surveys [19], student classwork [15], and whole class discussions

[7, 14]. The four common narratives that are pertinent to this study are technocracy, free market idealism, technological neutrality, and technological determinism.

Technocracy is a decision making approach guided by the belief that all problems can and should be solved through technology [21]. Here, technological solutions are privileged, often with little consideration of the social, political and historical context of the problems. Technocratic perspectives center the use of scientific and technological innovation to solve large-scale problems, such as world hunger and climate change, and decenter the underlying causes of these issues, which include human activity and politics [17]. By overemphasizing the use of technology to solve problems, technocratic framings remove from the solution space non-technical solutions that may include social and political action [17].

Free market idealism follows from neoliberal ideologies and emphasizes the supposed self-correcting nature of a free market that inevitably provides equal opportunities to everyone [20]. In this narrative, technological innovation, fueled solely by a free market, leads to progress, and hence leads to the benefit of all consumers [17]. However, in reality, technological development is guided by the values and needs of those with power in society and is both impacted by society and impacts society. As a result, those with power tend to benefit the most from technology as their needs are prioritized, and those with historically less power are more likely to be harmed [16].

The idea that technological development is always beneficial reinforces the common narrative of **technological neutrality,** the belief that technology itself is inherently objective and neutral in nature and any harm caused by technology is only manifested in how people choose to use that technology [20]. Technological neutrality ignores how social, economic, and political values are embedded in the development of technology [18]. These values tend to reflect the interests and needs of those who have historically held power in society [16, 20, 22].

The idea that technology is not neutral and is designed by humans also brings to question deterministic narratives of technology. **Technological determinism** assumes that technology develops in a self propelling fashion, where new technology is inevitable and humans must simply accept and adapt to it [23]. This framing absolves the creators of any responsibility in anticipating harm and designing to mitigate the negative impacts of the technology they design [17]. As technology is created by people, who live and work in societies, it inherently embodies the social norms, ideologies, and practices of societies [16, 20].

Since these narratives are pervasive, first-year engineering students are likely to have encountered them, and possibly accepted them, before they enter engineering programs. There is therefore a need to create classroom experiences that explicitly challenge these narratives, to develop students' sociotechnical understanding (or "critical sociotechnical literacy," as McGowan and Bell have proposed for K-12 students) [18]. As students are given opportunities to question these narratives, they will be able to recognize that the effects of technology are unevenly felt across groups of people and more-than-human actors and that who/what benefits and who/what is harmed typically aligns with historical power imbalances [16, 17, 18, 20]. These differential impacts result from the fact that the dominant social, political, and economic values in a society are reinscribed in technology, therefore reproducing and often exacerbating existing

inequities. As students see the connections between the design of technology and systemic injustice in societies, they will notice how who gets to define what counts as technology and engineering is coupled with who has power [18, 20]. By uncovering how technologies are situated within social and historical contexts, students can understand how the common narratives serve to benefit those in power – how when "technologies pose as objective, scientific or progressive, [they] reinforce racism and other forms of inequity" [16, p. 2].

Methods

Context

This study takes place at a medium-sized private university in the northeast United States. The context for this study is a first-year computation course for engineering students that is in the process of being redesigned to integrate sociotechnical content. There are five sections of the course, all taught by different instructors, with considerable autonomy, but all agreed to use project-created materials for the sociotechnical pieces of the course. At the time this data was collected, year 1 of the 3-year NSF-sponsored project, the first twenty minutes of one of two weekly course meetings was allotted for stand-alone sociotechnical discussions. To prepare for each discussion, students learned about a new sociotechnical case through an accessible reading or video and wrote a short reflection.

The discussions in this study come from week 11 of 14 of the semester. The topic of the week was algorithmic bias, specifically considering examples in Machine Learning and Artificial Intelligence (AI). The pre-assignment included two pieces focused on the work of Dr. Joy Buolamwini, self-described "poet of code": a video of a talk she gave titled, "Compassion through Computation: Fighting Algorithmic Bias" [24] and a comic highlighting her work created by NPR [25]. Both pieces focus on the consequences of algorithms and how those consequences are unevenly felt across different groups of people. Buolamwini describes how harm resulting from algorithmic bias does not have to be intentional, because the algorithms tend to reflect inequities already present in the world. She argues that there is a strong need for oversight, and also describes opportunities to combat computational bias, including creating inclusive data sets.

Participants & Data collection

This study focuses on first-year engineering students at a Predominantly White Institution. The university is making progress in becoming more diverse; the class year from which most of the students in this course belong is 58% female or gender minority, and 56% percent of the overall engineering student body is non-white [26]. A total of 192 students were enrolled across the five sections of the course in the target semester; of these, 84 were in the two sections of the course that were chosen for video recording. We video recorded one focus group in each of these two sections during each weekly sociotechnical discussion. The focus groups were chosen by practical considerations: of those groups composed solely of consenting students, we chose the group that was easiest to video record based on their location in the room. Since groups were determined by proximity–students worked with neighbors–there was some shifting of group composition over the semester. However, by week 11 of the semester, students sat in the same

seats and thus the small groups were quite stable. Table 1 provides pseudonyms and self-reported demographic information for the students in the two focus small groups.

Table 1: Pseudonym and demographic information for the students in focus groups

Pseudonym	**Gender**	**Self-reported race/ethnicity**	**Class year**	**Major**
Group 1				
Evan	Man	White	4th year*	Electrical Engineering
Laurel	Woman	White	1st year	Human Factors Engineering
Marina	Woman	East Asian, Hispanic, Latinx	1st year	Human Factors Engineering
Marisol	Woman	Southeast Asian	1st year	Mechanical Engineering
Group 2				
Ebo	Man	African (international student)	1st year	Electrical Engineering
Malik	Man	African American	1st year	Civil Engineering
Morton	Man	African American, Caribbean	1st year	Mechanical Engineering

*Although the course is intended for first year students, it is optional for certain majors, who are able to take it in later years if they desire.

Researcher Positionality

We identify as middle-class women who are interested in supporting and studying justice-oriented engineering education from elementary through undergraduate levels. Both of us have engineering degrees followed by engineering education degrees and have taught undergraduate engineering courses. The first author, Chelsea Andrews, is a white American early-career faculty member, has been a part of this research project from its inception, and leads the research component, including overseeing data collection. She designed many of the sociotechnical activities for the course, including the lesson analyzed in this study. Dr. Andrews is a mentor to the second author, Fatima Rahman, an international Ph.D. student in engineering education, who joined the project after video data for this study was already collected. The authors have a history of working together (as postdoctoral associate and masters student) on a research project in elementary engineering education.

Analysis

To begin to analyze small group sociotechnical discussions in engineering courses, we take a qualitative case study approach [27]. We choose one part of a single discussion, with data collected from two small groups, as we are interested in deeply understanding and characterizing individual student arguments [28]. We draw on discourse analytic techniques to analyze student talk, looking in cycles across grain sizes from individual phrases to turns of talk to the entire conversation. We focus on the ideas students express and look to understand their "meaning in social, cultural, and political terms" [29, p. ix].

In this exploratory analysis, we analyze both focus groups' discussions in response to the same prompt about using AI in hiring. This question was considered for 4 turns of talk in group 1 and for 3 turns of talk in group 2. To answer our research questions, we first identified arguments in students' discourse – looking for times when students made a claim and provided justifications (or warrants) of the claim [30]. Students' arguments were often developed over an entire turn of talk, which consisted of many phrases. The arguments were often accompanied by considerable hedging–many students began with phrases such as "I feel like" and "I think," and included qualifiers, such as "I don't know." Both authors identified individual student arguments separately and then cross-checked; there were no disputes. The authors then worked together to group similar student arguments; we found four groupings, presented in the discussion, but we hesitate to label these 'themes' as we were working with so few turns of talk. Finally, we drew on the literature on common narratives about technology to consider how these students' arguments aligned with or countered these narratives. The authors took each individual argument in turn and discussed which of the common narratives most connected to the arguments and claims and in what ways. Throughout the process, video recordings, transcripts, and identified arguments were shared within multiple engineering education research group meetings to hear different interpretations and increase the reliability of our findings.

Findings

Here we present and summarize the discussions of both focus groups in response to the question, "Do you think using AI throughout the hiring process is needed? What are the potential impacts of integrating AI into the hiring process?"

As you read these transcript excerpts from the two groups, notice how group 1 uses more general and neutral language, whereas group 2 is more explicit and uses more personal examples.

Group 1

For the first three minutes of their discussion, Group 1 members all agreed they liked the assigned comic and video and talked in fairly general terms about the negative effects of technology on marginalized groups. For example, Evan stated that, "tech has always been used to oppress marginalized communities in some way, shape, form or fashion so I am not surprised that this is happening again," and Marisol stated a similar understanding: "I feel it kind of reflects just like how the society has like kind of treated marginalized communities (.) and how (.) like the distribution of just like technology."

After this general discussion, Evan then shifts the discussion to the second prompt, whether AI should be used in the hiring process:

Line	*Speaker*	*Transcript*
1.1	Evan	I feel like it kind of falls into the hiring process (Laurel: yeah) yeah, and to your point [points thumb at Laurel] exactly like AI should never be used in the hiring process. I think that's kind of a easy quest–answer to this question of, just like it doesn't make sense. Because like removing the human to human interaction kind of keeps the employer from knowing the hire and

		also vice versa it keeps the hire from knowing their employer and so that kind of creates a level of distrust and (Laurel: yeah) yeah

Evan states that AI should not be used in hiring, using strong language such as "never" and "doesn't make sense." His reasoning is that it "creates a level of distrust" between the employer and the potential hire as AI would remove the contact between them. Evan is likely referencing a frame of the comic that showed an interview being conducted by an avatar on a computer screen. In this justification, there seems to be an implicit assumption that face to face hiring is mutually beneficial for both parties. It is unclear whether Evan is arguing that AI would do a worse job at hiring compared to humans (in terms of not hiring the same people) or that it would just be a worse experience for the humans and create a poor beginning to the working relationship.

1.2	Laurel	Although I think it would be very difficult to sort of ban it completely from the hiring process and so I was thinking sort of (.) there just needs to be very strict regulations for how it's used (.) and also a lot of transparency for how it's used. Um, because I don't think, you know, realistically you can make companies who are getting like, you know, so many applications and this would make their job, like I think, I don't know, I feel like it is sort of inevitable but we can make it at least more equitable. I think it's also like some things like AI as a whole is kinda hard to regulate but I think something as specific as hiring, I think that's like possible to regulate? I don't know, what do you guys think.

In her response to Evan, Laurel argues for regulation of AI in hiring. Her first phrase "Although, I think it would be very difficult to sort of ban it completely…so," appears to be in response to Evan's criticism, and can be read as 'I agree in principle, but think that's not possible, so here is a potential compromise.' She believes banning AI is not realistic, that the technology is "sort of inevitable," especially for companies who must process many applications. Laurel instead proposes regulations and transparency to "make it at least more equitable."

1.3	Marina	It can be used in like, not only like facial recognition, but like keyword searching through resumes (Laurel: right) where I feel like sometimes (.) like we've like gone through the college application process of like how people have to be like, 'oh it's a holistic view' where like, they'll just search for resumes, like jobs not colleges (Laurel: yeah), but like the resumes of- just like, let me see what GPA they had in college, or let me see if they have these skills and then not look at other things, where sometimes those other things are important. (Laurel: yeah) Where it is hard to do like hiring processes with so many applicants (.) without AI, but I don't think it's like <u>needed</u>. Like it's possible, it is harder, but sufficient, but like **note: phrases in quotes refer to times the speaker adopted an affected tone, which we interpreted to mean she was invoking a phrase spoken by others*

Marina brings up more details of how AI is or might be used in hiring. She mentions that AI includes not just technology such as facial recognition (which was mentioned in the video), but also keyword searching. She relates the hiring process to the college application process, where people claim "oh it's a holistic view," but in reality look at GPA and specific skills, and may miss other important parts of a resume or application. While Marina seems to be claiming that looking for specific markers is already happening–either manually or automated–her argument seems to imply that AI will continue this trend (while an argument could certainly be made that AI will make this worse, it is not clear if Marina is claiming this). Her conclusion is that hiring, even with "so many" applicants, is harder, but it is still possible, and thus AI is not needed in hiring.

1.4	Laurel	Yeah. I mean the consequences are huge, too, like it almost has the risk to sort of keep us stuck in–where we are. Since it's all like retroactive data, like, keeping us in the same loop of like, whatever current levels of diversity we have in the workplace. Um, yeah, there's a huge risk.

Laurel then raises the idea of consequences and risks of using AI, naming that it will keep us "stuck in–where we are…whatever current levels of diversity we have in the workplace". She recognizes that AI typically uses current hiring data in its decision making algorithms about what to look for in future applicants. As a result, there is a risk that future hires will be similar to those already hired.

This is the last turn of group 1 on this topic, before they pause and move on to whether anyone had any questions about the reading. Neither Laurel nor any others in this group explicitly mention which groups tend to benefit and who tends to be harmed by the AI bias, instead they use neutral language such as "current levels of diversity." As we will see in the next excerpt, Group 2 is explicit about who benefits and who is harmed by these technologies.

Group 2

Group 2 also began addressing the question of AI in hiring about 3 minutes into the discussion. Morton is the first to address the AI prompt:

2.1	Morton	So the second one is basically 'Do you think AI can be used throughout the hiring process and what are the potential impacts of it?' I said it doesn't need to be throughout the entire hiring process, because we already know that AIs are flawed based on everything we have gone through and the computing in the world* homeworks (.) so we see that AIs are flawed so if people end up – the hiring process building off the majority of people get hired – end up being white (.) And, as a result those who are like economically challenged, racially challenged–they won't be able to get a job because of the system. So if you need time for like, maybe a small portion of the job that doesn't need all these identity markers, it makes sense (unclear) provide (unclear) identity markers **"Computing in the world" is how we referred to this portion of the course*

Morton begins by arguing that AI does not need to be used "throughout the hiring process" because there is so much evidence that AI is flawed and builds off the majority. He explicitly states that more white people would be hired as a result and that people who are "economically and racially challenged" would be less likely to be hired. He ends with a caveat that some "small portion" of jobs might be reasonable to hire with AI if they do not need "identity markers"--but it is unclear what kinds of jobs he might be referring to.

2.2	Ebo	Um, what I said was, yeah, it doesn't need to be – like, it's helpful 'cause it's fast, like, you get your feedback as quickly as possible – but, there should be someone, there should either be a review–a constant review of this AI so, someone should actually like be doing their work today, that's what I feel, 'cause, I don't know if this is true but someone told me um, so when they are scanning your resume and all that they are just looking for keywords, like (Morton: Yeah) yeah I know that, so what if, I didn't know that. So, I think my first two or three times like applying for a job I just used like normal, random words – so assuming I didn't use those keywords, and actually have the qualification, because I don't – I didn't have those keywords in my resume, I'd not be hired. And then someone else who knows 'Oh, they have to add this to that' where they make it look–or package it and all that, compared to me, yeah. So I think, yeah that's what – you need someone who can actually see the whole, 'yeah this is not just words put together, but this is someone who can actually do the job', compared to someone who knows how to package themselves up. I think, yeah, that's all. That's what I said. I dunno

Ebo begins by recognizing the benefit of AI, that it is "helpful [be]cause it's fast." He then advocates for a constant review of AI, similar to Laurel (turn 1.2) in group 1 suggesting regulations. Ebo then discusses the role of keywords in applications, although it is unclear whether he's talking about humans scanning for keywords or AI. He invokes a personal story of applying for jobs two or three times before he learned that employers were looking for specific keywords in applications. Ebo relates that while he did have the qualifications, because he did not know to use the keywords, he was at a disadvantage compared to "someone who knows how to package themselves."

2.3	Malik	I said something similar. It definitely depends a lot on time. AI is probably the fastest way to like separate qualified and unqualified candidates, so many people just like default to that, but we like know the data has the biases of the people who like make the (unclear, possibly 'code'). Like sometimes, it will like find like uncommon names–like African sounding names–weigh them less than [air quotes] whiter sounding names. So it's kind of difficult. It's like not – like ideally it would be like, a group of people like analyzing the data itself or like you know, but it's kind of difficult to keep track. When you like look at the scale of large companies.

Malik begins by clearly stating one use of AI (and other processes) in hiring – separating candidates into qualified and unqualified bins and the need to focus on that part of the process. He then clearly states that the technology includes the biases that already exist, although it is unclear when he says "of the people who are like" if he is referring to the biased data that would result from building AI off of current employees or whether he is referencing the biases of the designers. Malik, like his group members, then gives another specific example of a problem with AI: that some software he is aware of distinguishes African sounding names and gives those names a lower rating than whiter sounding names. Finally, Malik acknowledges that this is a hard problem, because of the scale of hiring.

At this point, Morton shifts the group to the next question.

Discussion

While a technocratic stance – believing (explicitly or implicitly) that technology is the best solution to problems – is common throughout U.S. culture, including within engineering programs [3], we did not see any students clearly aligning with that stance in this data. While students stated their positions with a range of emphases, from "AI should never be used in the hiring process" (Evan, 1.1) to "it doesn't need to be throughout the entire hiring process" (Morton, 2.1), no student in either group argued that AI *should* be used in hiring. This is consistent with what we observed in discussions in the rest of the semester as well: students were generally skeptical of the technologies considered, in alignment with the readings and video resources that were assigned.

We are interested in these non-technocratic arguments that students stated–what reasons do students give for not using or not trusting technology? We saw students argue that (1) AI technology does not solve the problem well, (2) it is important to regulate AI, (3) using AI for hiring will stagnate diversity, and (4) using AI for hiring unfairly privileges some groups of people over others. Below, we describe each argument and analyze what common narratives the students are pushing back against by using that argument.

A note on the missing AI-human comparison: There are two common arguments for using AI instead of humans for hiring: (a) that humans are biased, while computers do not see physical characteristics, so must not be biased, and (b) that machines can do this process more efficiently, at a much larger scale than humans can do manually. In both groups' discussions, we see nods to reasoning (b) (turns 1.2 and 2.2), but neither of the recorded discussions nor any of these seven students' written pre-responses mention hiring done by humans. However, both groups mention examples, like the practice of looking for keywords, that could be accomplished by a human scanning resumes by eye, by a simple computer search function, or by a sophisticated AI program that is not necessarily looking for keywords inputted by a human, but is choosing words and phrases from a dataset of "successful" resumes. Therefore, it is unclear whether students (1) are not considering how hiring done by humans is similarly biased, or (2) recognize that both humans and machines are biased and are focusing on the problematic approaches to hiring, whether done by humans or AI. Because students did not explicitly compare human and AI-executed hiring, we are curious if this limits their analysis, but this is not a focus of this study.

Student argument: AI technology is not a good solution to the hiring problem

In group 1, there are a few justifications given for not using AI in hiring that center on how well the technology performs. Evan (1.1) argues against using AI because it will remove the "human to human interaction" and will keep the employer and hire from knowing each other well, which "creates a level of distrust." Marina (1.3) claims that AI would likely only be looking for information that can be searched, like keywords or students' GPA, and would miss other important aspects of an applicant. Both students are claiming that AI technology is not a good solution to the challenge that hiring poses for employers.

In their arguments, these students are resisting the common technocratic narrative that solutions that involve technology are inherently better than solutions that rely on human effort.

Student argument: It is important to regulate AI

One student in both groups focused on ways to make the technology safer. Laurel (1.2) argued for "very strict regulations" and "a lot of transparency" in AI. Ebo (2.2) said "there should be someone, there should either be a review - a constant review of this AI." In both cases, it seems like the students expect that this technology is going to be used in this way, but that there are potential safeguards to make it more equitable, even if they are not able to fully articulate what those safeguards would be. Laurel in particular seems to think that "realistically," because AI would make hiring so much easier, the technology is unavoidable: she states it is "very difficult to sort of ban it completely" and "I feel like it is sort of inevitable."

In arguing for oversight and regulations, these students are pushing back on the free-market idealism narrative. They are claiming that this technology is likely to cause harm, and that we cannot rely on the market alone to reduce this harm, but instead must have regulations. The narrative of free-market idealism would claim that technology will improve on its own in response to market pressures, in alignment with a technological determinism view that more technology is inevitable and the only solution to the problems created by technology is more technology. The students instead argue that the free market on its own cannot be trusted and in order to reduce harm we must look to non-technical solutions.

Student argument: Using AI for hiring will stagnate diversity

In group 1, only Laurel invokes a justification that considers the potential wider-scale impacts of AI technology in hiring. She states, "it almost has the risk to sort of keep us stuck in–where we are. Since it's all like retroactive data, like, keeping us in the same loop of like, whatever current levels of diversity we have in the workplace" (1.4). Laurel is implicitly claiming that the current levels of diversity are inequitable, and that the way AI technology would be used, relying on existing hiring practices and data, would simply replicate the status quo.

In this turn, by noting explicitly how the algorithm uses retroactive data, and that this leads to negative consequences, Laurel is rejecting both the technological neutrality narrative and the technological determinism narrative. Because the algorithms use data from society, it is clear here that influence between society and technology goes both ways, and that the technology itself, as it encodes these biases, is not objective but inherently harmful.

Since Laurel is considering the consequences of technology, and noting how the feedback loop of these algorithms would limit the diversity of the workplace, we can see that she is recognizing that applicants who are not the norm would be less likely to be hired. However, because Laurel only mentions "levels of diversity…in the workplace," her argument centers the workplace, rather than the applicants who are less likely to be hired by AI.

Student argument: Using AI for hiring unfairly privileges some groups of people over others

In group 2, all three students argue that AI would benefit some groups of people and harm others. Morton (2.1) states explicitly that "the majority of people get hired – end up being white (.) And, as a result those that are economically challenged, racially challenged–they won't be able to get a job because of the system." Ebo (2.2) notes that the hiring practice of looking for keywords privileges those who know "how to package themselves up" and disadvantages others, even if they have the qualifications. Malik (2.3) gives the example of technology giving a lower weight to "uncommon names–like African sounding names" compared to "whiter sounding names."

All of these students' arguments go against the common narratives of technological neutrality and technological determinism. They recognize that technology tends to benefit and harm different groups, and that these patterns tend to follow historical imbalances of power. They certainly do not believe that technology is inherently neutral and objective–they recognize that it encodes the biases of the culture, based on the biases of those who create it and the data it is based on.

Comparison of the two groups: specific language and centering harm

While both groups argued against using the technology and rejected the common narratives about technology, we see clear differences in the two discussions. The first difference is in language: group 1 tended to use generic language when talking about groups impacted by the technology, whereas group 2 is specific. The second difference is that group 2 consistently centers the people who are harmed in their discussion. We note that group 2 was a group of all Black males, whereas group 1 included white students and students with Latina and Asian ethnicities, which we believe influences how comfortable the two groups feel in having these conversations.

Taking group 1's conversation as a whole, it is clear they care about equity and justice and recognize that technology tends to reproduce injustice. They begin with observations that, "tech has always been used to oppress marginalized communities in some way, shape, form, or fashion so I am not surprised that this is happening again [Evan]" and are consistently concerned that the tech is not equitable. However, their discussion consistently uses only careful, safe language–the same language that universities and many academics use–like diversity and marginalized communities. Even Laurel, the only student in this group to explicitly raise an argument related to differential impacts of the technology, couches her reasoning in general terms such as "current levels of diversity." This group never explicitly mentions who in particular is harmed or that it is white people, and cis straight white men in particular, who have held power in society and continue to benefit, including from new technology.

In contrast, group 2 was explicit about who would benefit (white people) and who would be harmed by this technology (those who are economically and racially challenged, those who don't know how to package themselves, and those with uncommon names, including African-sounding names). They use specific, and sometimes personal, examples of how technology creates this harm. They recognize that those who benefit and those who are harmed follow historical power structures, and see how these power imbalances are reproduced through technology.

Rejecting common narratives without centering justice

While both groups reject the common narratives around technology, only group 2 focuses on who is harmed and who tends to benefit from technology. We can see in comparing the groups that group 2 consistently centers harm in their discussion: who specifically is harmed and how that harm is perpetuated through technology. Only group 2 prioritizes the "perspectives and welfare of the most vulnerable and marginalized" [17]. This difference between the groups is not evident from our original framework of focusing on the common narratives around technology– an unexpected finding of this analysis is that it is possible to reject the common technology narratives without centering justice concerns.

We conjecture that rejecting the common narratives of technocracy, free market idealism, technological neutrality, and technological determinism, is a necessary first step but is not sufficient for comprehending the complex interplay between technology and society. To fully understand how technology and society intersect, students must come to recognize: (1) that technology tends to have differential impacts, and that those who benefit and those who are harmed typically aligns with historical power imbalances; (2) that this is a result of the dominant social, political, and economic values of a society, which are embedded in technologies; and (3) how power and differential impacts directly relates to who gets to define what counts as technology and engineering in society [16, 18, 20, 22].

Conclusion

This study analyzed small group sociotechnical discussions in a first-year engineering course on computation and focused on students' arguments against using AI technology in hiring. We considered students' arguments through the lens of common narratives around technology: technocracy, free market idealism, technological neutrality, and technological determinism. Across the two student groups there were four main arguments put forward against using AI in hiring: (1) AI technology does not solve the problem well, (2) it is important to regulate AI, (3) using AI for hiring will stagnate diversity, and (4) using AI for hiring unfairly privileges some groups of people over others. All four of these arguments rejected the common narratives around technology. Overall, this finding aligned with our field notes from the classes and the reports from instructors.

This close analysis of the discourse data also revealed an unexpected nuance in the discussions: while students in both groups rejected the common narratives, only group 2 explicitly centered those who are harmed and how this harm would likely occur, and this group did so consistently. Group 1 managed to consistently reject the narratives while using vague, safe language, such as "current levels of diversity." They never explicitly mentioned who is harmed by the technology. Group 2, in contrast, explicitly states who will benefit and who is harmed. As a result, only the second group's discussion is clearly centered in justice concerns.

In future work, we will investigate how to scaffold small group sociotechnical discussions, what instructors should attend to during these discussions, and how to support students to orient toward systemic impacts. We are particularly interested in what pedagogical strategies would support students in having small group discussions that center justice. Creating and studying these supports will first require us to identify what repertoires of knowledge students need to adopt an justice lens in critiquing technology.

References

[1] Leydens, J. A., & Lucena, J. C. (2017). *Engineering justice: Transforming engineering education and practice.* John Wiley & Sons.

[2] Malazita, J. W., & Resetar, K. (2019). Infrastructures of abstraction: how computer science education produces anti-political subjects. *Digital Creativity*, *30*(4), 300-312.

[3] Cech, E. A. (2013). The (mis) framing of social justice: Why ideologies of depoliticization and meritocracy hinder engineers' ability to think about social injustices. *Engineering education for social justice: Critical explorations and opportunities*, 67-84.

[4] Pawley, A. (2019). Opinion: Asking questions, we walk: How should engineering education address equity, the climate crisis, and its own moral infrastructure?. *Advances in Engineering Education*, *25*.

[5] Riley, D. (2017). Rigor/Us: Building boundaries and disciplining diversity with standards of merit. *Engineering Studies*, *9*(3), 249-265.

[6] Claussen, S., Tsai, J., Boll, A., Blacklock, J., & Johnson, K. (2019, January). Pain and gain: Barriers and opportunities for integrating sociotechnical thinking into diverse engineering courses. In *Proceedings of the 2019 ASEE Annual Conference and Exposition*.

[7] Gupta, A. (2017, June). A practitioner account of integrating macro-ethics discussion in an engineering design class. In *2017 ASEE Annual Conference & Exposition*.

[8] Hess, J. L., & Fore, G. (2018). A systematic literature review of US engineering ethics interventions. *Science and engineering ethics*, *24*, 551-583.

[9] Winiecki, D., & Salzman, N. (2019, January). Analyzing and Working-Out Ways of Addressing Problems of Social-Justice in an Engineering or Computer-Science Context. In *2019 NSF REDCON (Revolutionizing Engineering & Computer Science Department CONference), Arlington, VA*.

[10] Gupta, A., Turpen, C., Philip, T., & Elby, A. (2019). Narrative co-construction of stances towards engineers' work in socio-technical contexts. *Critical, transdisciplinary and embodied approaches in STEM education*, 251-272.

[11] Ozkan, D., & Andrews, C. (2022, August), Perspectives of Seven Minoritized Students in a First-Year Course Redesign toward Sociotechnical Engineering Education *2022 ASEE Annual Conference & Exposition Proceedings*, Minneapolis, MN. https://peer.asee.org/41382

[12] Pangan, T. J., & Andrews, C. (2022, August), (Work in Progress) Examining how students critically evaluate racial bias in a medical device in a first-year computing course. *2022 ASEE Annual Conference & Exposition Proceedings*, Minneapolis, MN. https://peer.asee.org/41794

[13] Radoff, J., & Turpen, C., & Abdurrahman, F., & Chen, D., & Tomblin, D., & Agrawal, A., & Chudamani, S. (2022, August), Examining the "narrow" and "expansive" socio-technical imaginaries influencing college students' collaborative reasoning about a

design scenario *Paper presented at 2022 ASEE Annual Conference & Exposition, Minneapolis, MN*. https://peer.asee.org/41422
[14] Philip, T. M., Gupta, A., Elby, A., & Turpen, C. (2018). Why Ideology Matters for Learning: A Case of Ideological Convergence in an Engineering Ethics Classroom Discussion on Drone Warfare. *Journal of the Learning Sciences*, *27*(2), 183–223.
[15] Andrade, N., & Tomblin, D. (2019). What Are They Talking About? Depth of Engineering Student Sociotechnical Thinking in a Technical Engineering Course. *2019 ASEE Annual Conference & Exposition Proceedings*, 33551. https://doi.org/10.18260/1-2--33551
[16] Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code.* Polity.
[17] Gunckel, K. L., & Tolbert, S. (2018). The imperative to move toward a dimension of care in engineering education. *Journal of Research in Science Teaching*, *55*(7), 938–961. https://doi.org/10.1002/tea.21458
[18] McGowan, V. C., & Bell, P. (2020). Engineering Education as the Development of Critical Sociotechnical Literacy. *Science & Education*, *29*(4), 981–1005. https://doi.org/10.1007/s11191-020-00151-5
[19] Cech, E. A. (2014). Culture of Disengagement in Engineering Education? *Science, Technology, & Human Values*, 39(1), 42–72. https://doi.org/10.1177/0162243913504305
[20] Riley, D. (2008). *Engineering and Social Justice*. Springer International Publishing. https://doi.org/10.1007/978-3-031-79940-2
[21] Fischer, F. (1990). *Technocracy and the politics of expertise*. Newbury Park, CA: Sage.
[22] Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
[23] Smith, M. R., & Marx, L. (1994). *Does technology drive history?: The dilemma of technological determinism.* MIT Press.
[24] Buolamwini, J. (2019, March). Compassion through Computation: Fighting Algorithmic Bias [Video]. Talk given at the World Economic Forum, available at https://www.youtube.com/watch?v=_sgji-Bladk
[25] Cala, C., V. Stollberger, & L. Johnson. (2022, March). COMIC: How a computer scientist fights bias in algorithms. NPR. Retrieved from https://www.npr.org/2022/03/14/1085160422/computer-science-inequality-bias-algorithms-technology
[26] Freeman, L. R., Butler, C., & Sharkness, J. (Eds.). (2022). *Tufts University Fact Book 2021-2022*. Tufts University Office of Institutional Research. https://provost.tufts.edu/institutionalresearch/files/Fact-Book-2021-22.pdf
[27] Merriam, S. B. (1998). *Qualitative research and case study applications in education. Revised and Expanded*. Jossey-Bass, San Francisco, CA.
[28] Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative inquiry*, *12*(2), 219-245.
[29] Gee, J. P., (2011). *How to do discourse analysis: A Toolkit.* Oxon: Routledge.
[30] Toulmin, S. (1958). *The uses of argument*. Cambridge, MA: Cambridge University Press.