

Characterizing First-Year Engineering Students' Priorities and Language Use in Socio-technical Written Reflections

Dr. Kaylla Cantilina, Tufts University

Kaylla is a Postdoctoral Scholar at Tufts University where her work is motivated by design as a means for social justice. Her research explores the ways that students and practitioners seek to achieve equity in their design practices

Dr. Chelsea Joy Andrews, Tufts University

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

Fatima Rahman, Tufts Center for Engineering Education and Outreach

STEM Education graduate student at Tufts University

Context Matters: Characterizing First-Year Engineering Students' Shifting Perspectives when Reflecting on Different Sociotechnical and Justice Topics

Introduction

This practice paper relates to the overall mission of ECSJ, specifically transformative learning toward action. For well over a decade, there have been numerous critiques of the social/technical dualism present in engineering education and calls to disrupt it [1]. Researchers have varying motivations for this focus, including to better prepare students for engineering practice, which is inherently sociotechnical [2]; to increase the sense of belonging of historically excluded students, who are more likely to be interested in the social aspects [3]; and to create better societal outcomes that consider justice [4,5,6]. Attempts to disrupt the social/technical dualism and the apolitical nature of traditional engineering education have included revising stand-alone ethics courses and adding sociotechnical components to traditional engineering courses, such as design courses [7-10]. However, revising stand-alone ethics courses implicitly upholds the disconnect between the “technical” and “social,” and adding one or two modules to a traditional technical engineering course can be perceived as an additional load or “tack on” by both instructors and students. Furthermore, in addition to these structural barriers, sociotechnical content does not by default include a justice perspective. As such, an intentional focus on justice is necessary when designing pedagogical changes toward more holistic engineering education.

In response to these challenges and needs, we were funded by the NSF to conduct a research and pedagogical project in which we are integrating justice components throughout a first-year computing for engineers course. Instead of revising an ethics course or tacking on sociotechnical content to a traditional course, we chose to embed justice into our redesign of the “technical” class as much as possible. The revised course includes: (1) a weekly sociotechnical lab with small-group activities and discussions on curriculum-aligned real-world justice topics, (2) weekly post-lab readings and written reflections, (3) week-long projects where coding assignments are embedded in a justice topic, and (4) a final project that explicitly considers social impacts of numerical analysis or design. Each course section is supported by undergraduate equity learning assistants who help facilitate the sociotechnical labs and act as approachable peer mentors who can push students to think differently. Every assignment, including these written reflections, is graded and contributes to students’ overall course assessment.

While there were several artifacts, topics, and pieces of data regarding the class and its outcomes, this paper only focuses on analyzing how the weekly post-lab reading and written reflections illuminate themes in students' attention to and perception of the sociotechnical aspects of engineering. Specifically, we focus on reflections from one section of the five sections taught

during one semester. Our analysis aimed to address two research questions, the first regarding how students conceptualize bias, differential impacts, and the cause of societal outcomes, and the second regarding how these conceptualizations manifest differently across the weekly topics. In the future, we hope to do a more systematic inquiry into the rest of the data and other course components. Additionally, we hope that some of the findings in this paper might provide insight into our course design (specifically the design of the weekly reflections), student thinking, and potential future directions for both research and practice. These insights could prove useful to those at other institutions who are interested in using reflection activities in their own courses to center sociotechnical and justice components.

Methods and Analysis

Context

The context for this work is a first-year engineering computing course at Tufts University, a medium-sized private university in the northeast U.S. The course previously focused on teaching basic programming skills in Python or MATLAB, but has been re-designed to center sociotechnical and justice-oriented content throughout in-class small group activities and discussions, homework, and projects. This data comes from the second year of the study when the course included two “lecture” class meetings each week (although these were mostly active learning) and one “lab” meeting, where students engaged in small group activities and discussion on sociotechnical topics. While there were several other parts of the course that were redesigned, for this paper, we chose to only focus on the activity after each lab, which was a reading assigned along with reflection responses for students to complete individually. The topics and reflection questions for seven of the eleven weeks are listed below in Table 1. There were four weeks we did not include due to the data being so different that drawing comparisons between those and the other weeks was not possible: week 0 was an introductory week and the reflection questions did not center a specific technology or dataset, and weeks 5, 8, and 9 were project weeks with responses submitted by small groups (our unit of analysis was individual students). To ground the course content in social justice and sociotechnical principles that students could refer to throughout the course, we began the semester by introducing students to the course “sociotechnical tenets.” These tenets were referenced throughout the semester and occasionally used within reflection question prompts. The tenets are as follows:

- **Technology is not inevitable and does not always improve society.** Engineering and computing are subjective, require judgment, and require tolerating uncertainty – there is never one right solution or decision.
- **Data, algorithms, and technology are neither neutral nor objective.** Technology embodies the dominant values of society and the creators who design it, which tends to reproduce and/or exacerbate existing inequalities.

- **The effects of technology are unevenly felt across groups of people and more-than-human actors.** Who/what benefits and who/what is harmed typically aligns with historical power imbalances.

Table 1 below shows the week number in the left-most column, the topic of that week, and the reflection questions the students were prompted to answer after participating in their sociotechnical lab and readings.

Table 1. Descriptions of Weekly Topic and Assigned Reflection Questions

Week	Topic	Reflection Questions
1	Gender Differences in Car Crash Data	• How do gender inequities in car crash tests relate to the three sociotechnical tenets outlined in the first week?
2	Bias in Data Science	• What is bias? • How can definitions of bias influence data science? • What is the responsibility of the engineering team in these scenarios?
3	COVID-19 Modeling	• When making decisions, how should we use models? Do we use multiple models? • Do we research the models we are using beforehand? • What are the signs of a good model before we just trust the model that comes out first? When making decisions, how should we use models?
4	Designing for Disability	• How did these articles make you think differently about disability and design? • Why does Sasha Constanza-Chock believe it is important for communities to be “at the table”? Do you agree? • The NYT article states, “Disabled people have long been integral to design processes, though we’re frequently viewed as “inspiration” rather than active participants.” What does it mean to be an active participant in design? • What guiding questions would you suggest to engineering teams as they work with and for people different from themselves?
6	Impacts of Offshore Wind Implementation	• How might you relate this work to the sociotechnical tenets? • What questions do you have about offshore wind?
7	Race/ethnicity Disparities in Solar Deployment	• In both solar deployment and environmental racism, we see from the data that even after we account for income, race/ethnicity still matters in who is affected. Were you aware of this (race/ethnicity matters beyond income) before this week? Why do you think you were or were not aware? • What potential solutions can you imagine for either an environmental justice issue from the readings or the disparity in rooftop solar?
10	AI Technology in Hiring Processes	• Do you think using AI during the hiring process can be overall a benefit? In what circumstances? • What are potential impacts of integrating AI into the hiring process? • The metric right now for whether we should use AI is essentially, “is it less biased than humans”. Does that seem like a reasonable metric to use? What else could we use in place or in addition to that?

		• Why do most people with experience in AI not (yet) trust AI in hiring processes? What reasons do they give?
--	--	---

Participants & Data Collection

For this paper, we look closely at the written responses from students in one of the five sections of the course. Of the 180 total students who were enrolled in this course in this semester, 33 were in the focus section, and 25 of these students consented to participate in research on their coursework. We chose this section because it had both the most diverse set of students demographically and longer written responses, allowing us more insight into students' thinking. In future work, we intend to apply and adapt this coding scheme to the remaining sections. The chosen section consisted of two sophomores and 23 freshmen, including six white women, eight women of color, two women who did not report race/ethnicity, seven white men, one man of color, and one white nonbinary man. Students' race/ethnicity and gender were self-reported. Students' written responses were submitted via Canvas. After the semester ended and grades were submitted, the anonymized written responses were added to a spreadsheet.

Research Questions

Through this work, we sought to explore and answer two research questions regarding the nature of students' reflection responses.

1. How do students conceptualize bias, differential impact, and the cause of societal outcomes in weekly reflection responses focused on different topics?
2. How do these conceptualizations manifest differently across weekly topics?

Analysis

Our primary analytical approach was both inductive and deductive thematic analysis in alignment with qualitative research practices from Maxwell and Creswell [11], [12]. After putting the text data into a spreadsheet, our initial analysis began with summer undergraduate researchers immersing themselves in the data from the entire course, with the purpose of identifying interesting potential themes to investigate further. At this point, we had 12 themes, and the themes that were only found in one week were eliminated. After this first round of exploratory inductive analysis, three researchers (the co-authors) took a first pass at identifying themes loosely based on the themes distilled by the undergraduate research assistants. This first coding pass formed our code categories shown in Table 2 (cause or reason for situation, differential impacts, bias) which we then used in a second round of coding with a deductive approach to analyze the full dataset. After this, we conducted a final round of coding to inductively add nuance to the code categories which are described in Table 2. The unit of coding was each phrase, with sentences and responses able to contain multiple codes, or simultaneous codes including ones that logically conflicted with each other. For example, the bias code category has the subcodes "bias is not possible" and "bias is always present." A student might start their reflection with a phrase that describes bias as not being possible but then discuss bias

being present later in their reflection. The sentences coded for bias could also be coded as differential impacts etc. All the responses contained at least one code, though most of the responses did not have a code in each category. While the three coders did not code the entire dataset independently, at least two coders coded each week with a third coder checking the codes for the weeks they did not code. The collaborative coding was done entirely in person so that coding conflicts and inconsistencies could be addressed as soon as they emerged.

Findings

RQ 1: “How do students conceptualize bias, differential impacts, and the cause of societal outcomes in weekly reflection responses?”

RQ1 was primarily addressed by the creation of our codebook, as the codes characterize the conceptualizations. Through our analysis, we distilled the student conceptualizations into three code categories: (1) reason or cause for the situation, (2) differential impacts, and (3) bias. Each code category was further broken down into two codes, for a total of six codes. We created these codes because we saw a spectrum of perspectives within each larger code category. These codes helped capture how students were making sense of the weekly lab topics. These codes are described in Table 2, but in summary, are as follows: when students discussed the cause or reason for a situation in their responses we characterized them as *systemic* or *agentic*. Each discussion point that referred to differential impacts was coded as being either *harm-focused* or *neutral*. Each mention of bias was coded as either claiming that it is *possible to completely remove bias* from a dataset, algorithm, or technology, or that *bias will always be present* (that it is impossible to be unbiased).

The codes in each category described above (i.e., *systemic* and *agentic* within the *cause* category, and *neutral* and *harm* in the *differential impacts* category) seem like opposing perspectives, but we do not treat any of them as inherently more correct aspects to focus on in a sociotechnical analysis. At the same time, we know that students are generally more familiar with identifying *agentic causes* of problems, and we hoped that over the course of the semester, with practice applying a justice lens, students would improve their ability to identify *systemic* causes. Similarly, students noticing *differential impacts* with a *neutral* tone is reasonable, particularly in engineering spaces, but as students take up the justice focus of the course, we hope that students begin to notice that *harm* is particularly experienced by minoritized groups. The *bias* coding category, however, is different, in that there is an ascribed value for each subcode. A key course objective was for students to understand that algorithms, data, and technology are inherently biased and that neutrality or objectivity is impossible.

Lastly, most student responses had co-occurring codes; very few responses only contained one code ($n = 12$) or none at all ($n = 4$). Seven of the twelve single-coded excerpts were from the offshore wind week. The definitions and example quotes for each code category and code are shown in Table 2.

Table 2. Codebook

Codes	Description	Example Excerpt(s)
Code Category: Cause or reason for situation		
Systemic	The cause or reason for a flawed outcome of technology is due to the influence of systemic inequities/injustices, or society's values at the time of their creation	“Systemic racism perpetuates environmental racism, as putting people of color into certain communities gives larger companies the opportunity to put these people in a vulnerable position. Companies will put machinery in these communities because they know that these minority groups do not have the power to fight back against the pollution/health impacts they will face.”
Agentic	The cause or reason for a flawed outcome of technology is because of the intentional or unintentional decisions and actions of specific individuals such as engineers, business owners, policy-makers, etc.	“The engineers behind the Hybrid III failed to take into account that women and men have significantly different body types making it crucial to have different testing dummies structured accordingly to provide the most safety for passengers in the cars. The engineers designing this car only focused on “dominant values of society”, which at the time resulted in a male dominated viewpoint. This failure to see outside the dominant narrative, allowed for unsafe vehicles to be made that put women in severe danger.”
Code Category: Differential impacts		
Harm-focused	The impact of inequitable outcomes through technology is framed as the extent to which minoritized individuals/more than human actors and communities are harmed or bear the primary cost of innovation	“Research shows the people most affected by pollution are poor and non-white people. For example, black Americans are affected more than Asians and more than white people by air pollution of fine particulate matter type no matter what the income levels are.” [example of harm on more-than-human actors]: “The structures destroy valuable habitats and displace the organisms that live there, both of which could be detrimental to the ecosystems in the ocean.”
Neutral	The impacts of technology on society privilege some groups over others; harm is not discussed or centered	“I believe that the effects of offshore wind technology would be felt unevenly across groups of people. It makes sense that states that are located by large bodies of water would benefit from the addition of offshore wind compared to their landlocked counterparts.”
Code Category: Bias		
No bias is possible	It is possible that data or technology can become unbiased through some kind of process.	“If there is an algorithm that completely or nearly eliminates bias, then the benefits could help recruiters hire more efficiently. It could help recruiters eliminate candidates who don't have the required skills to work in the position and give hiring managers a smaller pile of resumes to review. If the algorithm is written without bias, it can also help eliminate human prejudice in hiring.”
Bias is always present	Data or technology is biased due to biased data collection or underlying societal bias. This may be intentional or not	“Though the creators of this AI were most likely not intentionally trying to be racist, their bias shows through in their work as they picked the white person to be the “standard” and didn't think to make sure their technology was inclusive or applicable to people of color.”

RQ 2: How do conceptualizations of bias, differential impacts, and cause manifest differently across weekly topics?

Through our analysis, we found that there were differences in what concepts were discussed in students' reflections across weeks and between each student. Figure 1 illustrates the differences in themes discussed in the written reflections based on the topic of each week on the horizontal axis, with the frequency of coded responses on the vertical axis. The colors in the stacked bar chart represent each code category with yellow representing "bias," purple representing "differential impacts," and teal representing "cause of situation." Within the colors, there are two different shades denoting the codes within each code category. The frequency of codes across the weeks is much less important to our research questions than observing the distribution of the codes within each week's reflections.




Figure 1. Frequency of Codes Per Weekly Topic

There are a few notable observations based on what is shown in Figure 1. First, each weekly topic facilitated the appearance of some concept categories more than others. The car crash, *bias* in data science, and hiring weeks contained the most frequent instances of our codes. These weeks also proportionally had more discussion of *bias* compared to the other weeks. In the solar deployment week, no students discussed *bias*.

Students discussed *differential impacts* during most weeks except for the COVID-19 modeling week, where very few students discussed *differential impacts*. Within the *impacts* code category, the topics, hiring, and bias in data science had more students discuss *impact as neutral*, whereas in solar deployment, most students discussed *impacts in a harm focused way*. The car crash week had a near half and half-split of students discussing impact as neutral and harm focused.

For the coding category, *cause*, car crash, bias in data science, disability in design, and solar deployment prompted students to discuss *cause as agentic*, whereas the weeks of COVID-19 modeling, solar deployment, and hiring saw more discussion of *cause as systemic*.

Second, the way codes were discussed manifested in a variety of ways across the different weekly topics. For example, when students discussed *bias* in the car crash week they primarily focused on the *agency of engineers* and car designers, as described by a student below:

“This lack of safety precautions [car designs being biased toward male bodies], was not a form of logic/reasoning, and was instead based on the ideas of the creator. Therefore, this demonstrates that inequalities can become more intense due to the beliefs of the creator.” Student 112

Conversely, though the AI in hiring week also had many students describe *bias* being present, *bias* was discussed in relation to algorithms and data (*with a possibility of there being no bias*) versus the ability of an engineer to influence the outcomes of their designs.

“I think that AI is still biased, but a different kind of bias compared to human bias. I think it still holds bias from the data that it uses to make its decisions on who is and isn’t qualified. If the data provided was not biased, then I think that its judgments would also be less biased.” Student 114

The two excerpts above from students 112 and 114 are representative examples of a larger trend across students in these two weeks: in the car crash week, many students discussed bias in relation to what “designers” can do, while in the AI in hiring week, students were often focused on characteristics of the technology.

Discussion about *differential impacts* also looked very different across the different topics. For the *harm-focused* subcode, some weeks elicited responses where students were *vague about who was harmed*, or in what ways, while other weeks students named *specific minoritized groups* that are especially impacted by the technology in certain ways. One such week was solar deployment, where many students focused on environmental racism, such as Student 127, who wrote:

“...solar panels are a new technology and I’d like to believe that the US is past discriminating due to race, but I don’t think it is [...] In the US, people of color

never had the same starting position as white people. So it is reasonable to assume that Black communities have faced the environmental burden disproportionality to other communities. This is referenced in the article where toxic waste was relocated to only Black communities.” Student 127

In contrast, during the offshore wind week, many students discussed *harm-focused differential impacts* less so in reference to a specific group of people, but more broadly about how the technology would impact low-income households:

“People who do not have enough money to pay electricity bills and those without access to electricity would not be able to feel the effects of their tax money. Additionally, if these projects cause electricity prices to rise, it will negatively affect low-income households.” Student 129

Here we see that while discussing the *harm-focused differential impacts* that low-income households are burdened with, student 129’s reflection did not include explicit discussions of systemic injustice versus financial cost-benefit as much as they were in solar deployment. We noticed that compared to other weeks, students grappled with history and anti-Blackness, as showcased by student 127, in solar deployment slightly more so than in other weeks.

There were also differences in how students discussed *cause* across the weekly topics. For example, in the designing for disability week, students identified the cause of non-inclusive products or lack of credit being given to disabled designers as able-bodied designers leaving users out of the design process. One student describes this aptly:

“Often a lot of product designers use the disabled as “inspiration”, meaning hypothetically the product should help them, but don’t involve a disabled person in the design process. This process is flawed as the product may not even end up being the most useful to the disabled person as it is based on a mere assumption from an able-bodied person about what disabled people need.” Student 123

This is noticeably different than in the COVID-19 SIR modeling week, where students were focused on *cause* being the effect of policies, countries, or governments compared to individuals or engineers. For example, one student relates the uncertainty of models as a reason why a quicker government response would have been helpful during the early stages of the pandemic:

“When multiple models are considered there are more variables to test to give a wider range of outcomes but also more error and sometimes too many things to bring into consideration [...] In a case like this (COVID-19 as an unpredictable virus) I think the country should have started off with contact tracing and quarantine to get ahead of the curve as other countries had found success [...]” Student 122

As the findings demonstrate, there is variance and nuance in how students discuss the themes outlined in our codebook. The student reflections explained above are only a few examples of how students discussed *differential impacts*, *bias*, and *cause* in their own ways within the scope of the codes. However, they do illuminate that not only are there differences across weekly topics, but also between students themselves within each week.

Discussion

In revising this first-year computing course to center sociotechnical content, our goals were to make visible to students the social aspects of engineering work and to help students cultivate critical thinking and reflection around the assumptions and decisions engineers make. To achieve these goals we designed the weekly labs and discussions to increase student familiarity with analyzing data, algorithms, and/or technology through a sociotechnical-justice lens. Given that this kind of sociotechnical analysis is relevant and necessary to all data science and engineering contexts, we utilized a survey-course approach where each lab topic was distinctly different. We hoped that this would help students adopt a sociotechnical lens in a wide variety of contexts, focusing on different marginalized populations. This breadth-over-depth approach taken in this first-year computing course would ideally be balanced by later upper-level courses where students could dive deeply into particular topics. We expected that different topics would make more salient different aspects relating to *cause*, *impacts*, and *bias*. Subsequently, our findings are consistent with this approach, as we see a strong variability in the focus of students' reasoning across topics. This can be seen visually in Figure 1, which organizes our results temporally and does not reveal any consistent temporal trend of students' beliefs and conceptualizations of bias, differential impacts, and causes over the course of the semester.

While we do not see a trend in students' conceptualizations of *bias*, *differential impacts*, and *cause* over time, we do see a notable variation in students' attention to different aspects of the sociotechnical topics. We propose that these differences are largely due to three key reasons: 1) affordances of the content of the topic, including the designed activity and assigned readings 2) students' familiarity with the topic, and 3) the nature of the discussion prompts.

For the first conjectured reason, affordances of the content of the topic, some topics, such as solar deployment, and their associated activities and readings were more explicitly tied to *systemic* injustice through the centering of issues like racism. This could be why weeks like solar deployment had more students discuss *systemic causes* compared to a topic like offshore wind. In contrast, although the Designing for Disability week (Week 4) was focused on justice in design, the accompanying readings focused on how engineers should leverage their expertise and recognize the design competency of disabled individuals. This focus on individual people (engineers and individuals with disabilities) may have been what led students to focus more on *agentic* rather than *systemic* causes for non-accessible technologies. Both the topics themselves,

along with the activities and accompanying readings could lead to the differing ways students conceptualized and contended with sociotechnical justice analyses. Similarly, some topics, activities, and readings connected more to data science than others, which could explain why some weeks had significantly fewer discussions of *bias* certain codes than others.

Another reason why we might not see a trend over time is that some topics might resonate with students' previous knowledge and experiences more than others. For example, we noticed that within the written reflections during the hiring week (week 10), students used rationale based on a belief that employers must review hundreds of thousands of resumes for an open position when this is a highly unlikely scenario in industry. However, because first-year engineering students are primarily 18-19-year-olds from the United States, many most likely have not experienced a rigorous job/internship search and hiring process. This could mean their reflections on the topic may leverage much less contextual understanding compared to other weeks. One example of such topics is the week on car crash data (week 2) where, given the United States primary dependence on cars for transportation, there is a higher likelihood of students being familiar with driving. Similarly, during the COVID-19 pandemic, many students were taking classes online from their own homes, and in most cities, public transportation was either shut down completely or service was greatly reduced. This time period would make a topic like the MBTA week more difficult for students to relate to compared to post-COVID, where many students use the public transit system to go out socially, or even commute to class. Consequently, based on our findings, we believe that it is important to think about the experiences of the student body when determining which topics to include in the labs in addition to how well they highlight sociotechnical justice issues.

The third reason we propose students' responses varied considerably each week is because of the nature of the reflection prompts. We expected in writing these prompts that students would write a general reflection on the lab activities and reading, making sure to address the questions in the prompts during their responses. Instead, most students answered each reflection prompt in order and often did not give more insights. Even with fairly open prompts like, "What questions do you have about ____", most students would either provide a very specific question or not respond at all. Overall, the prompts were much more influential in students' responses than we expected. As a result, if an aspect of sociotechnical thinking was not specifically raised in the prompt, then it rarely showed up in students' responses.

While we generally did not expect to see clear evidence of change over time in students' attention to different ST aspects, the one area we hoped to see change over time was bias. We know from teaching this course in previous years that students tend to come into the course with a technological neutrality lens that bias in data, algorithms, or technology, can be viewed as a numerical error that can be corrected. Further, students often have a technocratic view that technology naturally improves over time and biases will inherently be eliminated as technology

improves. Throughout the course, we attempted to push on these ideas, giving numerous examples to show how data, algorithms, and technology, are not and should not be seen as objective and neutral. Students' responses show that they are still wrestling with this concept, for example, writing in one week that bias is inevitable, and in a later week writing that as AI improves, bias in AI will be eliminated. While we will continue to emphasize bias throughout the class, we also recognize that the ideologies that students come into college with are strong and widely held, and it is likely not reasonable to expect a clear shift over just a few months of a single semester.

Conclusion

In this paper, we sought to characterize and analyze sociotechnical reflection responses from one section of a first-year engineering computing course. The reflections were embedded as assignments following the weekly labs which each tackled a different sociotechnical justice issue or topic every session, and the readings that were meant to be completed after class. Our analysis of the reflections from seven weeks produced three themes or coding categories: *bias*, *differential impacts*, and *cause of societal outcomes* which each contained subcodes. Using this codebook, we characterized the written reflections for each week; our findings demonstrate that some topics elicited more discussion of certain codes over others. Further analysis within the same codes but across weeks showed that students discussed the subcodes in different ways. As expected, because of the range of topics, there was little to no change over time in student attention to *bias*, *differential impacts*, and *cause of societal outcomes* with instead, a wide variation over weeks. We suggest that these findings can be explained by three factors: 1) the topics themselves and the associated readings and activities, 2) students' familiarity with the topic based on prior knowledge and experiences, and 3) the nature of our reflection prompts.

Based on these findings and the three aforementioned explanations, we are currently making adjustments to the course to better meet our learning goals: students learning the importance and necessity of sociotechnical justice analyses in engineering and data science. After our analysis, the findings identified offshore wind as a topic and set of resources that were not helping students conceptualize *bias*, *differential impacts*, and *cause of societal outcomes* as richly as we hoped. As a response, we plan to replace offshore wind with a different topic. To contend with the issues of student lack of familiarity particularly prevalent in the AI in hiring week explained in the discussion section of this paper, we replaced it with a lab on the U.S. Census which is a civic tool students are impacted by. We also thought more about how the topics could be ordered in such a way to give students prior knowledge that would better prepare them for another topic. For example, we taught the U.S. Census lab before solar deployment which helped students engage in an activity in solar deployment that used census tracts and data. Despite these topic changes, it is worth noting that just because some of our topics were not a good fit for our students, does not mean that topic could be effective in another setting. For example, the AI in hiring topic might be effective at an institution with a higher population of nontraditional

students or in classes for masters students who have come back to school after time in industry. Lastly, we have been updating the reflection prompts to be less open-ended, more tied to the assigned readings, and better challenging students to engage in critical thinking. An example of this is shown in Table 3 which displays a previous and revised reflection prompt from week 2.

Table 3. Example of Old and Updated Reflection Prompt for the Car Crash Week

Old Version	Updated Version
How do gender inequities in car crash tests relate to the three sociotechnical tenets outlined in the first week?	The two reports took different perspectives: the first concluded that disparate outcomes were because women are more fragile, while the second concluded that disparate outcomes were because cars are designed for the average male body. What are the possible implications for designing car safety features for each conclusion? Based on the reading, what values would you say are embedded in the way the US does car safety testing and regulations?

Some recommendations we have for others interested in implementing their own course redesign toward centering sociotechnical and justice components would be to be thoughtful about the connections between the technical skill building and the weekly topics so they do not seem disconnected from each other, or out of place. Taking into account these nuances and tailoring the course design accordingly is crucial to its efficacy. Instructors must consider what can and needs to be centered within each part of the course content. Some elements of a course center skill building and there may be a good and necessary reason for why sociotechnical content is not a good fit. For example, in the case of our class, a skill-building element was learning about passing variable names to a function. Adding sociotechnical content during this element could inhibit both technical and sociotechnical learning goals. Alternatively, it is useful to critique each part of the course and question if there is a *good reason* why some elements do not have a sociotechnical component. Consequently, the elements *without* a good reason are opportunities to integrate sociotechnical justice content. Lastly, this endeavor required constant iteration and collaborative efforts between the research team, instructors, and student equity learning assistants. As we continue to do more research to evaluate the short and long-term outcomes of the course, we hope that our preliminary findings prove useful to others in their pursuit of transformative engineering education.

References

- [1] Faulkner, Wendy. 2000. "Dualism, Hierarchies and Gender in Engineering." *Social Studies of Science*. 30(5):759–92.
- [2] Leydens, J. A., & Lucena, J. C. (2017). *Engineering justice: Transforming engineering education and practice*. John Wiley & Sons.
- [3] Malazita, J. W., & Resetar, K. (2019). Infrastructures of abstraction: how computer science education produces anti-political subjects. *Digital Creativity*, 30(4), 300-312.

- [4] 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.
- [5] 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.
- [6] Riley, D. (2017). Rigor/Us: Building boundaries and disciplining diversity with standards of merit. *Engineering Studies*, 9(3), 249-265.
- [7] 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*.
- [8] Gupta, A. (2017, June). A practitioner account of integrating macro-ethics discussion in an engineering design class. In *2017 ASEE Annual Conference & Exposition*.
- [9] Hess, J. L., & Fore, G. (2018). A systematic literature review of US engineering ethics interventions. *Science and engineering ethics*, 24, 551-583.
- [10] 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*.
- [11] Maxwell, J. A., *Qualitative Research Design: An Interactive Approach: An Interactive Approach*. SAGE, 2013.
- [12] Creswell, J. W., *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*. SAGE, 2014.