



Principles Matter: Integrating an Ethics Intervention into a Computer Security Course

Justin Petelka
University of Washington
Seattle, Washington, USA

Franziska Roesner
University of Washington
Seattle, Washington, USA

Megan Finn
University of Washington
Seattle, Washington, USA

Katie Shilton
University of Maryland
College Park, Maryland, USA

ABSTRACT

There is increasing agreement that teaching students ethics in computer science (CS) is important, but there is little agreement about how to teach ethics, when to teach ethics, or even what ethics curricula should include. CS programs are experimenting with both stand-alone courses and approaches that integrate ethics throughout the computer science curriculum. Drawing from work in CS education and Science & Technology Studies, we designed an integrated and interdisciplinary ethics intervention to help computer security students identify where ethics and politics intersect with their technical field and encourage students to see themselves as practitioners of politics and ethics. Through analysis of student assignments, post-course surveys, and instructor reflections, we found that, while our intervention had benefits for students and instructors, it only weakly encouraged students to think of themselves as practitioners of ethics and politics. Students also struggled to confidently adjudicate ethical dilemmas given only a set of ethical principles. Finally, the ethical principles we gave students strongly shaped their analysis – for example, students were more likely to consider disparate impacts of technology on marginalized groups when directly prompted to do so. Our results suggest that integrated and inter-disciplinary approaches have many benefits, but they require additional resources beyond a single course to effectively support students in adjudicating ethical dilemmas.

CCS CONCEPTS

• **Social and professional topics** → *Model curricula.*

KEYWORDS

ethics, computer security, social impact, assignments, university, undergraduate, ethical principles

ACM Reference Format:

Justin Petelka, Megan Finn, Franziska Roesner, and Katie Shilton. 2022. Principles Matter: Integrating an Ethics Intervention into a Computer Security Course. In *Proceedings of the 53rd ACM Technical Symposium on Computer*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCSE 2022, March 3–5, 2022, Providence, RI, USA.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9070-5/22/03...\$15.00

<https://doi.org/10.1145/3478431.3499275>

Science Education V. 1 (SIGCSE 2022), March 3–5, 2022, Providence, RI, USA.
ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3478431.3499275>

1 INTRODUCTION

There is increasing agreement that teaching students ethics is an important component of computer science (CS) education. CS programs are experimenting with both stand-alone [12] and integrated (“ethics throughout the curriculum”) courses [11], and interdisciplinary approaches [5, 17]. However, questions about how to teach ethics, what ethics training should incorporate, or even if CS-based ethics training alone can accomplish what is needed to address the “ethics crisis” [16, 30] linger in CS education. This is particularly true of computer security, a field in which research ethics and social responsibility have both been important topics for decades [1, 13, 23, 26]. While integrated and interdisciplinary approaches have many benefits, they create tensions between adding ethics material and retaining core course material, which lead to operational questions of time, effort, and responsibility [11, 17]. Best practices for teaching ethics in CS courses are still emerging, creating a need for experimentation and interdisciplinary collaboration in existing CS courses [11].

To contribute to this conversation, we report our experiences piloting an interdisciplinary ethics intervention in a computer security course. Our intervention, which we call Ask-an-Ethicist, draws on CS education research that suggests using integrated and continuous approaches to teach ethics [11, 17, 34], while our pedagogical goals are motivated by conceptualizations of ethics as “ethics work” from Science & Technology Studies (STS) [19, 20, 36]. “Ethics work” emphasizes that in practice, being ethical is an ongoing and political sensemaking activity in everyday technical work rather than strict adherence to a set of moral principles.

In an effort to position students as practitioners of ethics without sacrificing core course content, Ask-an-Ethicist altered existing course materials. Early in the course, the first assignment asked students (in groups) to identify stakeholders and power relationships in a computer security technology scenario, then write 3 to 5 related ethical questions they would want to ask a technology ethicist. The course instructor solicited ethical questions a second time as an in-class activity. For the second assignment near the end of the course, we synthesized student questions into seven questions, and asked students to revisit and individually answer a peer-generated question. At the end of the course, we brought in a panel of ethicists to address the student-generated questions.

Since we believed students may have limited experience in identifying and answering ethical questions in a computer security context, we provided students with one of three sets of ethical principles. Each set of principles present a different relationship to computer security: (1) the Menlo Report [1], an adaptation of traditional U.S. research ethics tailored for network and computer security research; (2) the Feminist Manifest-No [9], a purposefully political framework tailored for data ethics; and (3) Nussbaum’s capabilities framework [28], a generalized human rights framework that centers the baseline “capabilities” people need for human flourishing.

To evaluate our intervention, we analyzed student assignments, conducted a post-course survey of student impressions, and gathered feedback from the course instructor. We evaluated this combination of data sources along three different dimensions. First, we offer our reflections on designing an integrated and continuous interdisciplinary ethics intervention. We then evaluate how well we achieved our STS-influenced pedagogical goals. Finally, we examine whether and how the different sets of ethical principles shaped assignment responses

Our results suggest that our lightweight and integrated approach was easy for instructors to implement, did not diminish coverage of core technical topics, and made ethical questions interesting and relevant to students. Our approach has since been integrated into subsequent versions of the same course and into a separate Capstone course. However, students struggled in confidently applying their ethical principles in their assignments, and in post-course survey responses expressed apprehension about their ability to adjudicate ethical dilemmas given only ethical principles. We also found that the ethical principles we assigned strongly shaped which stakeholders students included, particularly whether students considered disparate impacts of technologies on marginalized groups.

Our experiences suggest that, while integrated and interdisciplinary ethics curricula (like ours) have several potential benefits, lightweight integration may require more complex and in-depth training for students to be successful. These findings complicate calls for integrated and interdisciplinary approaches without sufficient foundational scaffolding in moral reasoning and decision-making, and point to the importance of providing resources to help students develop broader ethical experience across a program instead of within a single course.

2 RELATED WORK

2.1 Teaching ethics in computer science

Our paper builds on and contributes to work at the intersection of ethics and computer science (CS) education. Researchers within the computing education research community have offered a number of pedagogical tools for teaching ethics in CS beyond traditional lectures, such as case studies [14, 18], gamification and role-playing [4, 7, 33] and science and design fiction [5]. Researchers have also explored methods for embedding ethics into existing technical courses [6, 11, 17, 31, 34], and analyzed the content of standalone ethics courses [12].

While standalone classes that teach ethics outside of the context of learning technology skills are more common [22], integrated approaches have several benefits. Addressing ethics alongside core

technology skills can challenge beliefs that ethics are separate from the everyday practices of the engineering profession [8, 12, 15, 30, 34] and habituate students to thinking ethically throughout system design development [17]. While our work is inspired by embedded ethics pedagogies in CS, total redesign of large classes are both resource [11] and labor intensive [17]. In this way, our work is most similar to Fiesler et al. [11], who alter an existing programming class assignments to “bring up ethical concepts and current ethical dilemmas in tech” while balancing the practical resource constraints of an individual class. Borrowing from these best practices, our intervention altered existing course assignments in an effort to balance surfacing ethical considerations in the everyday work of computer security without sacrificing core course materials.

Integrated approaches to teaching ethics can also benefit from interdisciplinary perspectives. Grosz et al. [17] developed ethics modules designed by PhD students and postdoctoral fellows in Philosophy and embedded these modules at least once within fourteen different CS classes. While post-survey responses suggested an overwhelmingly positive reception from students, the authors also noted institutional, knowledge, financial, and political barriers to sustainably upscaling these modules into a cohesive program. Raji et al. [30] suggest that these political barriers require our attention, arguing that AI ethics syllabi reflect an “exclusionary pedagogy” where ethics are distilled for computational approaches without deeper engagement with humanistic social science (HSS). The displacement of humanistic epistemologies “entails a loss of values, assumptions and methods that are crucial in HSS: hermeneutical, interpretative, qualitative methodologies and a sustained reflection on emancipatory societal goals” [30, p. 522]. As a result, the authors argue, computer scientists are trained to view problems of ethics through computational lenses without broader engagement of the limitations of CS-based epistemic approaches. Together, these bodies of work suggest that interdisciplinary approaches to teaching ethics is of interest to CS educators, and that best practices are still developing.

2.2 Ethics work in Science & Technology Studies

Our pedagogical goals draw on perspectives from Science & Technology Studies (STS) that position ethics as not merely the enactment of an ethical or moral framework, but as “ethics work” that is happening all the time [19, 20, 25, 36]. “Ethics work” is an everyday practice that may or may not include grappling with codified guidelines. Paradigms such as “ethical work,” “being ethical,” and “ethigraphy” highlight “how people organize themselves as ethical in the absence of the ontological security that professional ethicists and some philosophers presume” [36]. These concepts are analytical categories that do not denote whether an activity is “good,” but signal that people are grappling with the morality of an activity. In Ziewitz’s words, “[W]hat counted as ethical was thus not settled by applying guidelines to cast normative judgements, but by routinely sorting out specific practices as ethical or not.”

We drew upon these STS-conceptions of ethics to create two pedagogical goals for our intervention:

- Encourage students to identify themselves as practitioners of politics and ethics [20, 36]

- Help students identify when ethics and politics are happening in technical work [8, 11]

3 COURSE OVERVIEW

One co-author was the instructor of a ten-week undergraduate-level computer security course. The course consisted of 138 students that were primarily computer science (CS) majors, taught virtually in the Fall of 2020¹. The class met three times per week for an hour lecture, with one additional section led by the Teaching Assistants. The course's aims were (1) to help students in developing a “security mindset”, while maintaining a consistent eye towards the social impacts of technology, and (2) to develop fundamental skills in threat modeling and knowledge of modern attacks and defenses in computer security and privacy. Students were given a number of hands-on exercises in topics including cryptography, web security, software security, threat modeling, smart home security, and authentication. Course assignments consisted of two group and one individual homework assignment (25% of total grade), three group lab assignments (45%), a final project (20%), and a participation requirement that included in-class activities (10%).

4 INTERVENTION DESIGN

This section outlines the structure of our ethics intervention, followed by a description of the ethical guidelines we provided to students.

Three authors, previously unaffiliated with the course but with research and teaching backgrounds in ethics and STS, altered existing course materials with the course instructor. In total, we adapted two of the three course assignments (Assignment 1 and Assignment 2²), included an in-class activity similar to Assignment 1, and led an ethics panel at the end of the course. The timing of these four activities was spread evenly throughout the course in an effort to create a continuous and in situ intervention [34]. Because the bulk of our data from students comes from Assignment 1 and 2, we describe these two assignments in detail.

For Assignment 1 at the start of the course, students formed groups and were assigned one of the three sets of ethical principles; the Menlo Report [1], the Feminist Manifest-No [9], and the Capabilities framework [28]. We then asked students to (a) provide a summary of a computer security technology of their choice, (b) identify stakeholders and power relationships involved in the technology, and (c) generate 3 to 5 ethical questions related to this scenario that they would want to ask an ethicist. We provided students a one-page summary of each framework using direct quotes from the original documents as much as possible, as well as links to external explanatory resources to avoid significantly adding to the students' reading load.

Students responded to Assignment 1 and the in-class activity with a wide breadth of questions about ethics and technology security. The course instructor distilled over 150 individual questions from students into 9 categories, each with 1-3 specific questions. Then, three of the authors further synthesized this list into seven concrete questions:

- (1) When, if ever, should a government be able to ban a technology or application? For example, under what circumstances should a government be able to mandate that app stores remove a specific application?
- (2) Who should be held responsible for problematic activities that occur on platforms (e.g., encrypted messaging platforms, social media platforms, Tor)?
- (3) Should university-based research on computer vision techniques that enable “deepfakes” be stopped or paused?
- (4) How should companies be held accountable when security breaches occur or privacy violations come to light?
- (5) Under what circumstances should a government require companies to provide or build in backdoor access to encrypted technologies for law enforcement purposes?
- (6) Should homeowners be allowed to set up cameras that record what happens in a public space visible from their property?
- (7) Should parents have a right to monitor their children's use of technology? Alternatively (or additionally), should employers have a right to monitor their employees' use of technology?

Assignment 2 (near the end of the course) asked students to select and answer one of these seven student-generated questions. The goal of asking students to generate and answer classmates' questions was to position students as arbiters of a relevant and practical ethical dilemma. Presenting ethics in relevant contexts is a common practice in CS courses that teach machine learning techniques [31], and can increase success rates and retention in CS courses [11, 27]. We also hoped that placing students in the position of ethical adjudicator would help students see themselves as practitioners of politics and ethics in computer security work.

4.1 Ethical Principles

In an effort to support students unfamiliar with identifying and articulating ethical questions, we provided students with a set of ethical principles to ground their responses. We assumed no experience with ethics since CS students at the University of Washington are not required to take an ethics course. Because ethical principles and theories highlight and foreground different dimensions of an ethical problem [21, 29], we believed assigning different ethical principles would lead to variations in student responses. As such, students were assigned one of three sets of ethical principles (below), each presenting a different relationship to computer security. We asked students to use the same set of ethical principles for both of their assignments to familiarize them with a single set of principles.

Menlo Report (2012). The Menlo Report [1] was written by a group of computer and network measurement researchers, lawyers with expertise in computing, computer science research funders, and research ethics experts. The Menlo Report attempts to reformulate U.S. research ethics principles articulated in the Belmont Report (i.e., beneficence, justice, respect for persons) to guide computing research. The authors apply these principles to computing research and add the principle of “respect for law and public interest” to incorporate consideration of potentially illegal activities that are sometimes undertaken in computer security research. Because this framework was developed by computing security researchers, we

¹While this course is usually taught in person, due to the COVID-19 pandemic, this iteration of the course was taught virtually with lectures and sections meeting over Zoom.

²We have provided the full contents of both [Assignment 1](#) and [Assignment 2](#) here

believed it to be the most widely available set of principles which was specifically written with security dilemmas in mind.

Capabilities Approach (2011). The Capabilities Approach was originally developed by Amartya Sen to address the ineffectiveness of standard economic measures for analyzing the well-being of a country's citizens. Sen argued that analysts needed to focus on what people are actually able to do to have a flourishing life, which he often describes as “freedom” [32]. Sen's Capability Approach was formulated by the philosopher Martha Nussbaum into a “Capability Theory”, including ten fundamental capabilities that form a baseline for the possibility of living a life of dignity [28]. Although not frequently used in the domain of computing, the Capabilities Approach has been influential in the space of global development as a way to assess interventions [24]. We thought that students would be attracted to the clear human rights principles, and hoped that they could connect these to the ways in which sociotechnical systems enable or complicate human flourishing.

Feminist Data Manifest-No (2019). The Feminist Data Manifest-No [9] was written by a collective of feminist technoscience scholars during a workshop on feminist data practices. The Manifest-No consists of ten guidelines that are formulated both as negative (we refuse to) and positive (we commit to) normative statements. The Manifest-No authors explain their commitment to multiple feminisms, including Latinx, queer, Indigenous, Black, and trans-, and explicitly relates data gathering and issues to intersectional feminist ideals. While the Manifest-No, with its focus on data and is not a perfect match for a class focused on computer security, it does focus on technology and presents clear connections to histories of inequity. We chose these guidelines because they assume that technologies can never be neutral and ask the reader to consider how technologies actively marginalize different groups of people.

5 DATA COLLECTION AND ANALYSIS

Once the course was complete, we asked students for consent to analyze their assignment responses, incentivized by a \$100 Amazon gift card drawing³. While we asked students to contribute one group (Assignment 1) and individual assignment (Assignment 2), we only analyzed Assignment 2 to avoid an individual student contributing an assignment on behalf of their group.

In addition to contributing their assignments, we also asked student participants to complete a post-course survey. We adapted the survey from Grosz et al. [17], asking students whether each assignment component made ethics relevant and interesting; if each component helped students think clearly about moral issues; and whether the component increased their interest in learning about moral issues. In addition to the four Likert scale questions, we solicited feedback on each component of our intervention (i.e., Assignment 1, Assignment 2, the in-class activity, and the ethics panel) through a free-response textbox. Out of 138 students, 26 (19%) students offered to contribute their Assignment 2 responses for analysis (Assignment Participants), and 20 (15%) students completed the survey (Survey Participants). We do not know why students did not participate - they may not have been reading instructor emails after the course ended or just not wanted to participate. We

suggest that researchers should never assume that students want to participate in research and that this warrants further research.

To analyze student assignments and survey free-text responses, we used an inductive thematic coding process [3] to identify themes present in student responses. We divided our coding process into two rounds to sensitize ourselves to differences between and within student assignments. For the first round of coding, three authors looked at all of the responses for one of the three ethical frameworks. During the second round, three authors looked at all of the responses to specific ethics questions. We then used the results from the survey to supplement or contextualize our findings.

6 RESULTS

Combining student assignments, survey data, and reflections from the instructor, this section describes the results of the intervention and our reflections. In summary, embedded approaches to teaching ethics in CS have many benefits but require broader support for students to confidently adjudicate ethical dilemmas. Our intervention also did not lead students to see themselves as practitioners of ethics and politics in technical work. Last, we found that the ethical principles that we provided students really mattered: students were more likely to consider impacts to marginalized communities when explicitly prompted to do so.

6.1 Benefits of embedded approaches

Post-course survey responses [Table 1] were generally enthusiastic about our assignments, with over 70% of students strongly agreeing that our assignments were relevant and interesting, helped them think about moral issues, and increased interest in learning about moral issues. We do note that self-selection to participate in the study might have also shaped our survey results. Students appreciated seeing the ethical guidelines twice (once during each assignment). Survey Participant 6 (SP6) suggested that “[l]ooking at the framework a second time helped me remember it and having the guiding question helped me see it in a new context/try to apply it”.

The course instructor/co-author found that having the space to formally articulate ethics questions and concerns seemed to benefit students. The instructor was impressed with the breadth of questions students submitted in Assignment 1. Since these questions were generated in the second week of the course, this suggests students were engaged and interested in questions of technology ethics before the class began. Our assignments may have served as an outlet for students to articulate and grapple with social concerns related to technology security. We note that these questions were mostly debates prominent in the news and not specific to computer security work (e.g., botnets or penetration testing). Collecting these questions later in the course when students have been exposed to the material may help surface questions specific to the work of computer security.

The instructor also found the embedded and course-long nature of the intervention to be a good fit since the class already engages themes of social concerns invoked by computer security throughout the course. In addition, adapting existing assignments minimized the time taken away from core course materials. The portability of our assignments allowed them to be integrated into both subsequent versions of the course and a separate Capstone course.

³This protocol was approved by our institutional IRB.

Question	Assignment 1				Assignment 2			
	was interesting	was relevant to me	helped me think more clearly about the moral issues we discussed	increased my interest in learning about the moral issues we discussed	was interesting	was relevant to me	helped me think more clearly about the moral issues we discussed	increased my interest in learning about the moral issues we discussed
Strongly agree	5	6	8	8	11	8	11	12
Agree	9	8	10	6	6	10	6	5
Somewhat agree	5	5	1	4	3	1	2	0
Neither agree nor disagree	0	1	1	1	0	1	0	3
Somewhat disagree	1	0	0	0	0	0	1	0
Disagree	0	0	0	1	0	0	0	0
Strongly disagree	0	0	0	0	0	0	0	0

Figure 1: Post-course survey responses for Assignment 1 and 2 of Ask-an-Ethicist

6.2 Struggles with ethical principles

Despite the enthusiasm shown in our Likert scale questions, students frequently articulated difficulties in using ethical principles in their free-text responses. Many students felt the guidelines were “shoehorned in”, leading them to feel like they were “jumping through hoops to forcibly connect the framework” to their question (SP1). The most common feedback item we received was for students to choose their ethical principles rather than have them assigned. This suggests that students found some frameworks to be more amenable to particular ethical questions than others.

Students also frequently requested scaffolding, examples, or explanations for “using” ethical principles. Asking for samples of this exercise, SP9 said “I was a bit nervous at the time [that] what I was submitting was not what the instructors were looking for”. Other students reflected on their own lack of experience in adjudicating ethical dilemmas: “[ethics are] absolutely something I feel like as a graduating senior I was not exposed to enough” (SP4). Our reading of student assignments indicated that their self-assessments were accurate – many students had difficulty in consistently applying the principles to assess and consider their ethical questions. Both survey responses and assignments suggest that students need more than ethical principles to confidently adjudicate ethical dilemmas.

6.3 Students did not see themselves as ethics practitioners

One of our goals was to help students see ethics as an everyday “practical accomplishment” [20] meaning that they themselves were involved in moral work and could be agents of change. We coded student responses for the stakeholders they identified to understand how they saw themselves in relation to the ethics prompts. Generally, the majority of responses identified only stakeholders explicitly named in their ethics question (e.g., government, technology companies, etc). No respondents positioned themselves as having agency or responsibility in their ethics scenario, or explicitly incorporated themselves into groups with responsibility such as “researchers” or “homeowners”. This suggests students did not directly envision themselves as practitioners of politics and ethics in technical work, but rather addressed the question from the position of the stakeholders provided. This could be an artifact of how we selected and synthesized the ethics questions from students. In retrospect, we should have re-framed student questions that

more obviously asked them to see themselves as ethical actors. We discuss this in more detail below.

6.4 Ethical principles matter

Our analysis showed that the ethical principles assigned had a large influence on student responses. While students generally did not consider stakeholders beyond those presented in the prompt, students using the Feminist Manifest-No were an exception. These students were much more likely to identify and discuss the consequences of an ethical dilemma for marginalized stakeholders (6/10 responses using the Manifest-No) compared to students using the Menlo Report (0/8 responses) or the Capabilities Approach (2/8 responses). This was the clearest theme that emerged from our coding. For example, Assignment Participant 9 (AP9) considers that, “because of the racial injustice that has happened in past as well as what is happening currently, people of certain color and race may be more vulnerable in investigations, even if innocent”.

The Manifest-No responses also described approaches to ethical deliberation not found in responses using Menlo and Capabilities guidelines, suggesting that students understood sociotechnical systems as actively producing social outcomes for different peoples. AP3 suggests self-reflection as a key component of ethical decision making; that homeowners should “review their own biases” before deciding to install a public-facing doorbell camera. Manifest-No responses also discussed including a broader variety of stakeholders in decision-making processes, such as when AP8 suggests that deepfake researchers “take into account the consequences of their research at every step and consult with marginalized communities”. Taken together, these results suggest the Manifest-No may support students in foregrounding and grappling with histories of discrimination and the situated concerns different people experience with technologies. If one of the goals of ethics teaching is to promote discussions about discrimination and structural inequity as a result of sociotechnical systems, interventions should explicitly require this consideration.

7 LIMITATIONS, DISCUSSION, AND FUTURE WORK

Our intervention was not designed as a controlled experiment and our sample size was small (26 assignments, 20 survey responses out of 158 students). However, consistent results across assignments,

survey responses, and instructor reflections gives us confidence in our interpretations.

Our intervention successfully made ethics in computer security interesting and relevant, and was easily adapted to existing course material and other courses. However, our activity was only marginally successful at helping students identify where ethics and politics intersect with their technical field and unsuccessful at encouraging students to see themselves as practitioners of politics and ethics in technical work.

Though students successfully constructed questions that identified ethics and politics in computer security, the questions generated by the class were largely focused on national-level debates prominent in the news (e.g., the Trump administration's TikTok ban, doorbell cameras), rather than topics over which students might have direct impact in their current work or near-future careers (e.g. developing botnets, penetration testing, bug bounties). The content of the questions in turn framed both the stakeholder groups considered. The actors in student responses were often broad - governments, corporations, researchers, and the general public - but never did a student respond with what they might do in a given situation. In retrospect, the questions that we chose from student respondents did not ask students to reflect on what they might do, or how they might have power in an ethical dilemma. Future adaptations to this assignment will explicitly position students as empowered ethical agents.

Our findings suggest that students were more likely to discuss disparate impacts of technology on marginalized communities when directly prompted to do so by the ethical guidelines (i.e., the Manifest-No). Consequently, when ethics curricula specifically aspire to have students discuss the impacts of technology on marginalized groups, instructors should embed this prompt directly. In our approach, this could happen either by selecting specific points of the Manifest-No that foreground differential vulnerabilities and histories of discrimination. An alternative approach is suggested by our finding that the actors named in an ethical question were powerful determinants of the stakeholders students considered in their responses. Offering ethical questions that highlight complicated social relationships - for instance, asking whether non-custodial parents rather than simply "parents" should be able monitor their children's technology use - might also increase students' engagement with vulnerability and non-normative relationships.

Finally, we struggled to create assignments that could supplement rather than replace existing technical content while providing enough scaffolding for students to be comfortable and confident in their ethical analyses. As we saw in the surveys, students often struggled to understand and use the ethical guidelines. While it may be possible to provide more scaffolding [5, 17], this would either take time away from core course materials or require external resources (e.g., a PhD student versed in ethics and teaching [17]). An alternative would be for CS students to take a foundational class in social sciences and humanities to train students in strategies for interpretation and hermeneutics, then incorporate embedded ethics interventions in multiple classes to reinforce and practice foundational skills [30]. This would have the added benefit of not overburdening individual courses [11].

Another way to provide additional clarity would be through a guided peer review assignment, giving students an opportunity to

refine their responses to an ethics question. While not a perfect solution, incorporating guided peer review has the potential to give students experience and confidence in articulating and critiquing ethics work, challenge student assumptions, and other pedagogical benefits [10].

8 CONCLUSION

The assignments that we designed and the responses that we received provoked reflection about our curriculum and how we teach ethics in computer science broadly. We gave students different ethics guidelines to provoke them to analyze technical situations from points of view other than their own - a task which was moderately successful, and more successful for guidelines which explicitly delineate attention to power and historical injustice.

In particular, we found that student responses were unlikely to articulate concerns for marginalized populations unless prompted to do so from the ethical guidelines. Helping students unpack and expand their imaginations of technology stakeholders is a critical component for humanistic values of emancipation and inclusive participation [2, 30]. Our findings suggest that if instructors want students to engage with concerns of bias, disparate impacts, and differential experiences with technology, instructors must make these concerns explicit.

Last, students struggled to confidently apply ethical guidelines. Taking a particular ethical approach and effectively using it to analyze a situation is a skill learned over time. The contextual application of ethical rules or guidance is a part of *phronêsis*: ethical wisdom, which must be developed through practice [35]. However, the practice of applying ethical frameworks is difficult to achieve outside of a stand-alone ethics course. Currently, there is considerable momentum towards incorporating ethics into computing education classes and curricula. Our study suggests that ethics interventions like ours can help CS students see ethical questions as interesting and relevant without taking time away from core course materials, and that explicit guidelines can guide reflection on diverse stakeholders and disparate impacts. However, it also suggests that lightweight interventions such as ours may struggle without more complex and in-depth training for students. These findings complicate calls for in situ approaches without sufficient foundational scaffolding in ethical frameworks and decision-making, and point to the importance of providing resources to help students develop broader ethical experience.

ACKNOWLEDGMENTS

We thank our ethics panelists, Quinn DuPont, Patricia Garcia, and Michael Zimmer for their help making this project a success. Additional thanks to all the students who participated in our intervention; Erika Wolfe and Eric Zeng for supporting classroom activities and discussions; Tadayoshi Kohno for providing feedback on the project and for his work developing earlier versions of the course and assignments; and Alannah Oleson for reviewing drafts of this paper. This work was funded by the National Science Foundation, Grant #1634202.

REFERENCES

- [1] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo Report. *IEEE Security Privacy* 10, 2 (March 2012), 71–75. <https://doi.org/>

- 10.1109/MSP.2012.52
- [2] Ruha Benjamin. 2019. *Race after Technology: Abolitionist Tools for the New Jim Code*. Polity, Cambridge.
 - [3] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
 - [4] Bo Brinkman and Keith W. Miller. 2017. The Code of Ethics Quiz Show. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education (SIGCSE '17)*. Association for Computing Machinery, New York, NY, USA, 679–680. <https://doi.org/10.1145/3017680.3017803>
 - [5] Emanuelle Burton, Judy Goldsmith, and Nicholas Mattei. 2018. How to Teach Computer Ethics through Science Fiction. *Commun. ACM* 61, 8 (July 2018), 54–64. <https://doi.org/10.1145/3154485>
 - [6] Mary Elaine Califf and Mary Goodwin. 2005. Effective Incorporation of Ethics into Courses That Focus on Programming. In *Proceedings of the 36th SIGCSE Technical Symposium on Computer Science Education (SIGCSE '05)*. Association for Computing Machinery, New York, NY, USA, 347–351. <https://doi.org/10.1145/1047344.1047464>
 - [7] Roxanne L. Canosa and Joan M. Lucas. 2008. Mock Trials and Role-Playing in Computer Ethics Courses. In *Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education (SIGCSE '08)*. Association for Computing Machinery, New York, NY, USA, 148–152. <https://doi.org/10.1145/1352135.1352187>
 - [8] Erin A. Cech. 2014. Culture of Disengagement in Engineering Education? *Science, Technology, & Human Values* 39, 1 (Jan. 2014), 42–72. <https://doi.org/10.1177/0162243913504305>
 - [9] M. Cifor, P. Garcia, T.L. Cowan, J. Rault, T. Sutherland, A. Chan, J. Rode, A.L. Hoffmann, N. Salehi, and L. Nakamura. 2019. Feminist Data Manifest-No. <https://www.manifestno.com>
 - [10] Maria De Marsico, Filippo Sciarone, Andrea Sterbini, and Marco Temperini. 2017. Supporting Mediated Peer-Evaluation to Grade Answers to Open-Ended Questions. *EURASIA Journal of Mathematics Science and Technology Education* (2017). <https://www.ejmste.com/article/supporting-mediated-peer-evaluation-to-grade-answers-to-open-ended-questions-4708>
 - [11] Casey Fiesler, Mikhaila Friske, Natalie Garrett, Felix Muzny, Jessie J. Smith, and Jason Zietz. 2021. Integrating Ethics into Introductory Programming Classes. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. ACM, Virtual Event USA, 1027–1033. <https://doi.org/10.1145/3408877.3432510>
 - [12] Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*. Association for Computing Machinery, New York, NY, USA, 289–295. <https://doi.org/10.1145/3328778.3366825>
 - [13] Megan Finn and Quinn DuPont. 2020. From Closed World Discourse to Digital Utopianism: The Changing Face of Responsible Computing at Computer Professionals for Social Responsibility (1981–1992). *Internet Histories* 4, 1 (Jan. 2020), 6–31. <https://doi.org/10.1080/24701475.2020.1725851>
 - [14] Kenneth R Fleischmann, Russell W. Robbins, and William A. Wallace. 2009. Designing Educational Cases for Intercultural Information Ethics: The Importance of Diversity, Perspectives, Values, and Pluralism. *Journal of Education for Library and Information Science* 50, 1 (Winter 2009) (2009), 12. <https://www.jstor.org/stable/40732558>
 - [15] Natalie Garrett, Nathan Beard, and Casey Fiesler. 2020. More Than “If Time Allows”: The Role of Ethics in AI Education. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*. Association for Computing Machinery, New York, NY, USA, 272–278. <https://doi.org/10.1145/3375627.3375868>
 - [16] Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. 2019. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. *Hawaii International Conference on Science Systems*, 2122–2131. <https://doi.org/10.24251/HICSS.2019.258>
 - [17] Barbara J. Grosz, David Gray Grant, Kate Vredenburg, 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. <https://doi.org/10.1145/3330794>
 - [18] Christina Harmon and Chuck Huff. 2000. Teaching Computer Ethics with Detailed Historical Cases: A Web Site with Cases and Instructional Support. *ACM SIGCAS Computers and Society* 30, 3 (Sept. 2000), 24–25. <https://doi.org/10.1145/572241.572251>
 - [19] Catherine Heeney. 2017. An “Ethical Moment” in Data Sharing. *Science, Technology, & Human Values* 42, 1 (Jan. 2017), 3–28. <https://doi.org/10.1093/hhrp>
 - [20] Klaus Hoeyer, Aaro Tupasela, and Malene Bøgehus Rasmussen. 2017. Ethics Policies and Ethics Work in Cross-National Genetic Research and Data Sharing: Flows, Nonflows, and Overflows. *Science, Technology, & Human Values* 42, 3 (May 2017), 381–404. <https://doi.org/10.1177/0162243916674321>
 - [21] Anna Lauren Hoffmann and Katherine Cross. 2021. *Teaching Data Ethics Foundations*. Technical Report. University of Washington. <https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/46921/TeachingDataEthicsFoundations-Hoffmann-Cross.pdf?sequence=5&isAllowed=y>
 - [22] Rick Homkes and Robert A. Strikwerda. 2009. Meeting the ABET Program Outcome for Issues and Responsibilities: An Evaluation of CS, IS, and IT Programs. In *Proceedings of the 10th ACM Conference on SIG-Information Technology Education (SIGITE '09)*. Association for Computing Machinery, New York, NY, USA, 133–137. <https://doi.org/10.1145/1631728.1631764>
 - [23] Maritza L. Johnson, Steven M. Bellovin, and Angelos D. Keromytis. 2012. Computer Security Research with Human Subjects: Risks, Benefits and Informed Consent. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, George Danezis, Sven Dietrich, and Kazuo Sako (Eds.). Springer, Berlin, Heidelberg, 131–137. https://doi.org/10.1007/978-3-642-29889-9_11
 - [24] Chad Kleist. 2010. Global Ethics: Capabilities Approach. <https://iep.utm.edu/ge-capab/>
 - [25] Michael Lynch. 2001. The Epistemology of Epistemics: Science and Technology Studies as an Emergent (Non)Discipline. <https://asaskat.files.wordpress.com/2015/11/skat-331.pdf>
 - [26] Kevin Macnish and Jeroen van der Ham. 2020. Ethics in Cybersecurity Research and Practice. *Technology in Society* 63 (Nov. 2020), 101382. <https://doi.org/10.1016/j.techsoc.2020.101382>
 - [27] Hilarie Nickerson, Catharine Brand, and Alexander Repenning. 2015. Grounding Computational Thinking Skill Acquisition Through Contextualized Instruction. In *Proceedings of the Eleventh Annual International Conference on International Computing Education Research (ICER '15)*. Association for Computing Machinery, New York, NY, USA, 207–216. <https://doi.org/10.1145/2787622.2787720>
 - [28] Martha C. Nussbaum. 2013. *Creating Capabilities: The Human Development Approach* (reprint edition ed.). Belknap Press: An Imprint of Harvard University Press, New Delhi.
 - [29] Michael J Quinn. 2020. *Ethics for the Information Age* (8th edition ed.). Pearson. <https://www.pearson.com/content/one-dot-com/one-dot-com/us/en/higher-education/program.html>
 - [30] Inioluwa Deborah Raji, Morgan Klaus Scheuerman, and Razvan Amironesei. 2021. You Can’t Sit With Us: Exclusionary Pedagogy in AI Ethics Education. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 515–525. <https://doi.org/10.1145/3442188.3445914>
 - [31] Jeffrey Saltz, Michael Skirpan, Casey Fiesler, Micha Gorelick, Tom Yeh, Robert Heckman, Neil Dewar, and Nathan Beard. 2019. Integrating Ethics within Machine Learning Courses. *ACM Transactions on Computing Education* 19, 4 (Aug. 2019), 32:1–32:26. <https://doi.org/10.1145/3341164>
 - [32] Amartya Sen. 2000. *Development as Freedom* (reprint edition ed.). Anchor, New York.
 - [33] Katie Shilton, Donal Heidenblad, Adam Porter, Susan Winter, and Mary Kendig. 2020. Role-Playing Computer Ethics: Designing and Evaluating the Privacy by Design (PbD) Simulation. *Science and Engineering Ethics* 26, 6 (Dec. 2020), 2911–2926. <https://doi.org/10.1007/s11948-020-00250-0>
 - [34] Michael Skirpan, Nathan Beard, Srinjita Bhaduri, Casey Fiesler, and Tom Yeh. 2018. Ethics Education in Context: A Case Study of Novel Ethics Activities for the CS Classroom. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18)*. Association for Computing Machinery, New York, NY, USA, 940–945. <https://doi.org/10.1145/3159450.3159573>
 - [35] Shannon Vallor. 2016. *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (1st edition ed.). Oxford University Press, New York, NY.
 - [36] Malte Ziewitz. 2019. Rethinking Gaming: The Ethical Work of Optimization in Web Search Engines. *Social Studies of Science* 49, 5 (Oct. 2019), 707–731. <https://doi.org/10.1177/0306312719865607>