What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis

Casey Fiesler
casey.fiesler@colorado.edu
University of Colorado Boulder
Boulder, CO

Natalie Garrett
natalie.garrett@colorado.edu
University of Colorado Boulder
Boulder, CO

Nathan Beard
nbeard@umd.edu
University of Maryland
College Park, MD

ABSTRACT

As issues of technology ethics become more pervasive in the media and public discussions, there is increasing interest in what role ethics should play in computing education. Not only are there more standalone ethics classes being offered at universities, but calls for greater integration of ethics across computer science curriculum mean that a growing number of CS instructors may be including ethics as part of their courses. To both describe current trends in computing ethics coursework and to provide guidance for further ethics inclusion in computing, we present an in-depth qualitative analysis of 115 syllabi from university technology ethics courses. Our analysis contributes a snapshot of the content and goals of tech ethics classes, and recommendations for how these might be integrated across a computing curriculum.

CCS CONCEPTS

• **Social and professional topics → Computing education**.

KEYWORDS

curriculum, ethics, professional responsibility, syllabi

ACM Reference Format:
Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *The 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20), March 11–14, 2020, Portland, OR, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3328778.3366825

1 INTRODUCTION

In the wake of scandals like Cambridge Analytica [9], Google employee protests over military contracts [32], and biased algorithms in Amazon's hiring processes [24], many people are wondering if computer science is facing an "ethics crisis" [42]. These examples are all from 2018, though since the inception of digital computing as a field, scholars have noted the potential for computing technologies to raise ethical and social issues that are fundamentally different than those relevant to other contexts [38]. Discussions of the role of ethics in computer science (CS) education also go back as far as the earliest SIGCSE conferences [26]. However, as scandals like these increase in frequency and technology ethics becomes more prominent in both public consciousness and scholarly attention [38], it is unsurprising that we have also seen an increase in related university courses— along with demand for even more [35].

Particularly given calls for greater integration of ethics content across the CS curriculum [16, 30, 36], it is likely that an increasing number of CS instructors will be including ethics as part of their courses. However, a potential barrier to this type of integration is that instructors without a background in the area might feel that they do not know what or how to teach when it comes to ethics [30]. There have been many descriptions of individual ethics-related pedagogical practices in the computer science education research community [4–6, 16, 36], though the wide variety of technology ethics courses available today across different disciplines also provide a window into broader patterns of the content and goals of education in this space.

Through a qualitative analysis of 115 syllabi for university technology ethics courses and an analysis of metadata from a total of 202 such courses, we present a description of a broad slice of technology ethics coursework. This paper contributes analyses of (1) in what departments and by whom tech ethics classes are being taught; and (2) what self-described "tech ethics" courses include for both content and learning outcomes. Taken together, this work contributes a snapshot of the content and goals of tech ethics classes and what role these might play in the broader CS curriculum.

2 BACKGROUND AND RELATED WORK

How to conduct ethical, just technology design; to avoid biased or otherwise harmful practices; and to consider the goals and politics of technology and whether the interactions or other consequences it creates are fair and good, are all parts of a longstanding conversation within computer science [34]. Of course, what constitutes "good" in this context is not a given, nor is what "ethical" might mean [34]. Beyond "ethics," there are a number of related concepts, such as values, justice, and responsibility, and conversations rooted in different traditions or literature (such as values in design or business ethics [15]). As Stark and Hoffmann point out, ethics (particularly in the context of professional codes) should be a starting point, not an end point, for the kind of world we want to see with respect to technology's outcomes [39]. However, in this paper, we will use the term "tech ethics" in order to examine what *is* being taught under this potentially broad designation.

Strategies for ethics education, in terms of both content and structure, has been a topic of conversation in many different disciplines for many years. For example, the fields of law [29], medicine [33], business [20], and engineering [2] have all considered the benefits of infusing ethics throughout an entire curriculum versus covering

SIGCSE '20, March 11–14, 2020, Portland, OR, USA

ACM ISBN 978-1-4503-6793-6/20/03…$15.00
https://doi.org/10.1145/3328778.3366825

relevant content in standalone classes. The benefits of applied ethics education that is discipline-specific are well-established—both in encouraging moral development generally [20] and convincing students that ethics is an important component of their profession [20, 29]. However, a number of studies have also shown that ethics instruction often does not translate into experiences outside the classroom [3, 13]. Therefore, there are still many questions about how to best accomplish the goals of ethics education, and the ways that different universities and programs teach ethics is far from homogeneous in both content and extensiveness [8].

How to best approach ethics education within computing is also a longstanding question; it first appeared at SIGCSE in 1972, where in a discussion of social responsibility and computing education, Nielson suggested that it is the responsibility of CS instructors to "try to prepare the student to make better decisions" [26]. In the U.S., the accreditation board ABET requires that CS programs must produce students that have "an understanding of professional, ethical, legal, security and social issues and responsibilities" but does not specify how this should be accomplished [30].

In 1996, a report from an NSF-funded project to inform CS ethics curriculum recommended that ethics content should be integrated into core CS classes as well as taught in standalone classes [23]. Integration across the curriculum acknowledges that ignoring ethical issues as they arise in-situ for other learning marginalizes ethics; it should be presented as a necessary part of daily practice rather than a public relations digression from what is really important [29]. However, this strategy is not common practice, and ethics content in technical CS classes (even those for topics for which ethics is highly relevant like machine learning) is rare [28]. Possible reasons for not integrating ethics into a technical class could include an instructor not knowing what to teach or how to teach it, or not thinking that it should be part of the material [30]. In response to a 1996 *Communications of the ACM* article that proposed ethics integration, a CS professor wrote a letter to the editor arguing that ethics is "not computer science" and that it was "difficult to imagine a computer scientist teaching these things" [22]. However, in recent years as technology ethics has become a more important part of public discourse, there has been an increase in attention to ethics in computing pedagogy, particularly for emerging technology topics such as data science and artificial intelligence [35].

Within the computing education research community, we have seen detailed descriptions of specific options for ethics pedagogy—for example, gamification [4, 6], immersive theater [37], incorporation of science fiction [5], integration into HCI [36] and machine learning [30] classes, and use of codes of ethics [31, 41]. However, as with other disciplines, there is not a common core curriculum in computing ethics. A study of ethics syllabi in medical curricula emphasized the role that this kind of data can play in aiding individual faculty members in developing their own ethics content, by providing a broader picture of patterns and possibilities in this space [8]; here, we hope to do the same.

Analysis of syllabi is a common method for considering curriculum requirements in educational research [10], and there have been efforts within the computing education research community to use syllabi to understand current strategies for teaching computer science [1, 40]. A course syllabus can be used to assess both the structure of a course and the exact knowledge units covered, as well as how individual learning objects are packaged for students [40]. In this paper we describe our analysis of over 100 syllabi, in order to provide this broader picture for technology ethics education that can help inform ethics-related instruction across the CS curriculum.

3 METHODS

For this study, we collected and qualitatively analyzed 115 syllabi from university courses self-described as "tech ethics." We informed this analysis with the broad goal to understand the content, character, and goals of tech ethics classes.

3.1 The Dataset

In November 2017, controversy around a New York Times op-ed that claimed academics were "asleep at the wheel" when it comes to tech ethics [27] inspired the creation of a crowdsourced collection of "tech ethics" syllabi in the form of an uncurated, openly editable Google spreadsheet [11]. Anyone could add a class, and the information collected includes course title, university, department, instructor, level (i.e., undergraduate or graduate), date last updated, and links to the course description and/or syllabus, though not all entries include all of this information. The spreadsheet was shared widely; within a few days of its creation, there were over 50 entries [7], and as of August 2019 there are over 230. At the time of our data collection in November 2018, there were 202.

Importantly, though the spreadsheet included links to publicly available information, it was created as a public resource, and not as part of a research project. Therefore, before analyzing this data, we contacted every instructor listed on the spreadsheet. For those who included links to their syllabi, we gave them the opportunity to opt out of having their syllabus included in our data; for those who had not included links, we asked if they would like to have their class included. As a result, our dataset includes only those syllabi that were publicly available and linked to from the spreadsheet or were explicitly provided to us, and we removed those where the instructor opted out of inclusion. The final dataset that we qualitatively analyzed includes 115 syllabi. However, we did conduct some aggregate analysis on the metadata included with the spreadsheet entries from the entire 202.

We do not make claims about the representativeness of this dataset. We conducted this analysis with the goal of providing examples of ethics instruction, with an eye towards trends and patterns. However, our data likely over-samples from instructors who are comfortable sharing their course materials and/or are active on social media. Our aggregate measures are also intended to provide a snapshot of tech ethics classes, not to make broad claims about the overall state of the field. We provide these aggregate measures in part to describe the data, so that findings can be interpreted with its character in mind.

3.2 Data Analysis

After collecting copies of these syllabi as they appeared in November 2018, we mined each syllabus for, when available, topics as listed in a schedule or as presented in a reading list, and learning outcomes or objectives. Not all syllabi included these components. We conducted two separate qualitative analyses in which two researchers worked together to create codebooks with thematic categories that could be applied to the entire dataset [21]. For the topic analysis our goal was to group the instructor-described topics into higher-level categories, so we used affinity diagramming to thematically group topics before conducting formal coding of all topics for each class [18]. For the learning outcome analysis we conducted iterative, open coding to construct a codebook before applying those categories to the entire dataset [21]. All analyses included at least two researchers during each stage, and during the analysis process we discussed edge cases and disagreements to come to consensus.

We also collected and analyzed the metadata from all 202 syllabi from the crowdsourced spreadsheet, in that it represented information about where the classes are being taught and by whom. To supplement the information provided, we used Google to find publicly available information about home departments and disciplines of named instructors where available. We used affinity diagramming to create higher-level categories for disciplines, mapped department and instructor disciplinary background to these categories, and then calculated descriptive statistics that we have incorporated into our findings.

4 FINDINGS

In describing the overall picture of tech ethics education as presented by this large sample of syllabi, we include descriptions of: (1) in what department and by whom these classes are being taught; (2) content taught, categorized by topics; and (3) learning outcomes in these classes.

4.1 Who teaches tech ethics?

Though the courses in our dataset are likely not a representative sample of all university tech ethics courses, they represent a cross-section of types of universities. The dataset contained 94 different universities (since a number of universities had multiple courses listed). Though the majority are in the U.S. (70 out of 94), there is also representation from Canada, Europe (6 in the UK, and also France, Switzerland, Austria, Portugal, Spain, and Denmark), South America, Australia/New Zealand, and the Middle East. The majority appear to be taught in English, with a few in Spanish. These courses also represent both private and public universities, and range from R1 research universities to liberal arts and teaching colleges.

In early discussions of computing ethics pedagogy, one prominent question was "who should teach it?" In 1994, Johnson argued that philosophers and social scientists trained in ethics should be teaching these courses so that they have the appropriate expertise, though a counterargument from Martin was that it is critical that regular CS faculty teach in order "to emphasize to students that social impact issues are a fundamental part of computer science, not some tangential topic that they take somewhere else" [19]. These are both reasonable arguments, and one solution has been to involve both philosophy and CS instructors in course design [16], though this is a resource-intensive solution that is not always possible.

Table 1 shows the disciplinary breakdown by home department of the class, home department of the instructor, and discipline of the instructor's terminal degree(s). We had metadata for 202 courses, though some courses were cross-listed, some courses had multiple instructors, some courses had multiple degrees, and some types of information were unavailable for some courses. As indicated in Table 1, the courses in our dataset are taught in a range of different types of departments, though the most common are computer science, information science, and philosophy. Regardless of the department, the class may or may not be targeted specifically at computer science majors, as tech ethics is also a topic that is relevant for many areas of study.

Table 1: The number of classes for which the course home department, instructor home department, and instructor degree matches each discipline, sorted by course home most to least.

Discipline	Course Home	Instructor Home	Degree
Computer Science	67	61	31
Info Science	62	49	36
Philosophy	26	21	40
Communication	23	18	19
Other Non Tech	18	18	20
Sci & Tech Studies	13	6	13
Engineering	12	10	7
Law	11	13	22
Other Tech	9	8	7
Math	7	3	6
Business	3	4	1

Interestingly, though most classes are taught within computer science, and the most common home department for instructors is also computer science, it is more common for the instructor's disciplinary background (as represented by their terminal degree) to be in philosophy or information science than in computer science. Note that though each course is represented once under "course home," some instructors taught multiple courses or had multiple degrees, and each instructor is only counted in "instructor home" and "degree" once. We also looked at the level of the class, when this information was available. 107 were taught at the undergraduate level, 74 at the graduate level, and 19 cross-listed for both; for the undergraduate classes, 67 appeared by the course number to be upper division (junior/senior) and 50 lower division (freshman/sophomore).

Our analysis also revealed that the faculty teaching the classes are fairly well-distributed between different faculty ranks—16%, 18%, and 16% for assistant, associate, and full respectively as the highest rates, with instructors, adjuncts, and graduate students making up the remainder.

4.2 What topics are part of tech ethics?

The majority of courses out of the 115 we qualitatively analyzed were general tech ethics courses that covered a range of topics, though 19 were context-specific. By far the most common specific topic for an entire class was artificial intelligence, with 10 AI-specific courses. There were also courses specific to disability, law, videogames, natural language processing, and bioengineering.

Our topic mapping was based on how the topic was described by the instructor. This means that, depending on how a class was organized, a particular sub-topic might appear in a different category;

Table 2: The number of courses that had content for each listed topic, out of 115 total courses, organized from most popular to least popular topic.

Topic	Courses
Law & policy	66
Privacy & surveillance	61
Philosophy	61
Inequality, justice & human rights	59
AI & algorithms	55
Social & environmental impact	50
Civic responsibility & misinformation	32
AI & robots	27
Business & economics	27
Professional ethics	25
Work & labor	23
Design	20
Cybersecurity	19
Research ethics	16
Medical/health	12

for example, one professor might put a discussion of GDPR under a class titled "regulations" and another might put it in a class titled "privacy." We mapped each topic to only one higher-level category, and in very rare cases when an overarching course topic clearly represented multiple categories (e.g., "privacy law"), we made a judgment call based on the description.

Our qualitative analysis of topics resulted in a set of 15 high-level categories, which we created through an iterative process of affinity diagramming. Though a number of topics appeared in the syllabi that did not fit into any of these categories, for the purposes of this description, we synthesized the analysis into topics with at least 10 instances in our dataset of 115 syllabi. These categories and their relative frequencies are detailed in Table 2. Here, we explain each category, with more detail for the most common topics.

The most common topic, appearing in 57% of the syllabi analyzed, was **law and policy**. This is unsurprising, as over half of the syllabi in our dataset combined ethics and policy as the overarching focus of the class, and a number of classes included applying or understanding law as a course objective. Courses that included law and policy content typically covered policies that specifically impact technology, like the General Data Protection Regulation (GDPR), the Digital Millennium Copyright Act (DMCA), and other legal concepts such as free speech and intellectual property as they apply to computing technology. Law and policy topics also intersect with other topics such as artificial intelligence, particularly in the context of case studies and news stories that discuss the legal implications of topics like predictive policing, algorithmic discrimination, or whistleblowing.

The next most common topic was **privacy and surveillance**, covered in 53% of the courses analyzed. This topic covers a lot of ground, from anonymization in the context of big data and social media to surveillance in the context of smart cities, social scoring, and wearable technology, and raises issues such as personal responsibility, digital footprints, and data as property. This topic also intersects with other domains, and is discussed in the context of a number of specific computing topics, such as research ethics, design, security, and data science.

Traditional ethical theories and other aspects of **philosophy** are covered in 53% of the syllabi analyzed. Many courses begin with an overview of one or more ethical theories, with the most common being utilitarianism, deontology, virtue ethics, and social contract theory. When courses in computer science departments covered ethical theory, utilitarianism was by far the most common, though we noted a wider variety of theories in classes housed in philosophy departments. We also see courses cover content around morality, cultural norms, and how to develop a personal code of ethics as a result of the course.

A little over half of the syllabi analyzed covered some aspects of **inequality, justice, and human rights**. Though a broad topic, courses that covered this category tend to highlight how technology can reproduce and augment existing societal discrimination, marginalization of populations, and inequality. This includes discussions of inclusivity and accessibility that determine who gains access to technology and is able to participate in the knowledge production that results, as well as inequality in how technology is designed.

Artificial intelligence (AI) as a general topic appeared often enough that during analysis we differentiated between AI in the context of algorithms and AI in the context of robotics and automation. **AI and algorithms** was more prominent, covering topics like black box algorithms and machine learning. While there is some overlap with the inequality and justice topic (and which category the topic was mapped to depended on how the instructor sliced it), the content of this category most often deals with algorithmic fairness, bias, and profiling. Many courses discuss the imperative of transparency and the challenge of auditing algorithms. 48% of the courses included this AI and algorithms, whereas **AI and robots** appeared in 23% of courses. This emerging topic in AI often covered wartime and labor applications of these technologies, as well as issues of explainability, regulation, and morality in the context of self-driving cars, drones, autonomous weapons, and intimacy.

Though the impact and consequences of technology is an overarching theme across computing ethics, some topics focused specifically on large-scale **social or environment impact**. For example, some courses discuss the paradox of technology increasing social connections while also diminishing in-person social structures. Environmental concepts include green computing, environmental degradation, sustainability, and geoengineering. We also included in this category topics related to the social impact of harassment that are afforded by technology, such as doxing, hate speech, revenge porn, and trolling. Closely related are the topics of **civic responsibility and misinformation**, which appear in 28% of syllabi. These courses include topics related to individual and collective responsibility for maintaining existing democratic values and combatting disinformation and propaganda online, as well as causes of polarization, such as filter bubbles, fake news, misinformation, societal manipulation, and propaganda.

Business and economics and the related topic of **work and labor** each appeared in a little less than a quarter of courses. Business and economics covered capitalism, financial models, marketing, pricing (ethics of freemium models, for example), advertising, and the free market. Many courses also cover case studies of ethical dilemmas and controversies at various corporations, and/or called for corporate social responsibility and responsible innovation. Courses that covered work and labor as a broader topic focus less on corporations and more on conceptual topics such as job and wage loss as a result of automation, outsourcing, or changing job markets. These topics highlighted solutions like universal basic incoming, crowdwork, and encouraging workforce diversity and organization.

Given that not all of the courses in our dataset are specific to computing majors (e.g., courses taught in Communication or Science and Technology Studies), it is not alarming that only 22% of courses included topics dedicated to **professional ethics**. Within those courses, risk analysis and responsibility to society were emphasized. Five classes included a specific reference to the ACM Code of Professional Ethics.

A handful of additional topics were covered respectively in less than 20% of the courses. **Design** as a topic included tangible ways to design technology that could have a positive impact on society, including discussions of design for accessibility, value-sensitive design, and design fiction. **Cybersecurity** included any topics related to hacking or cyber-attacks, cyber terrorism, or international cyber property. **Research ethics** covered the way researchers use information and what is appropriate, including the Belmont Report, informed consent, A/B testing, p-hacking, and the ethics of data scraping for research. **Medical/health** topics covered medical research ethics, healthcare, genetic engineering and enhancement, or bioethics.

Many of the topics listed here were the organizational structure for readings or other content.[1] The readings listed for these topics were largely a mixture of books, scholarly articles, and news stories. Textbooks were uncommon, and a number of courses also included media like documentaries, television shows (e.g., *Black Mirror*), and science fiction films and novels.

4.3 What are the goals of teaching tech ethics?

The terms "learning objectives" and "learning outcomes" have subtle differences in meaning but are often used interchangeably [17]; we saw both in our data. Similar to a 2019 syllabus analysis published at SIGCSE [1], here we use the term "learning outcome" to mean an explicit statement of what students should know or be able to do by the end of a course, regardless of how such a statement might be labeled in any given syllabus. Our qualitative coding revealed eight common types of learning outcomes in terms of what students should be able to do by the end of the course. Though some outcomes included specific content knowledge (e.g., "understand intellectual property issues" or "differentiate between ethical theories"), these largely mapped to the topics in the previous section, and were more rare. Of the 115 syllabi we analyzed, 100 included learning outcomes; Table 3 shows the breakdown of how many courses included each category of outcome.

Learning outcomes in these classes focused far more on conceptual skills than on specific knowledge–particularly compared to learning outcome analysis for other computer science classes such as intro programming [1]. Common keywords included verbs like "critique," "evaluate," "reflect," and "analyze." The difficulty of evaluating these kinds of learning outcomes (compared to, for example, code that can be evaluated on whether or not it compiles) is one challenge for CS instructors in teaching this content [25].

Table 3: The number of courses that had each type of learning outcome, organized from most common to most common outcome.

Outcome	Courses
Critique	71
Spot issues	36
Make arguments	26
Improve communication	26
See multiple perspectives	23
Create solutions	21
Consider consequences	18
Apply rules	10

The overarching goal of ethics courses appears to be to teach students to recognize ethical issues in the world (similar to the "issue spotting" skill that is critical to legal education [14]), critically evaluate and assess these issues and technologies (including considering multiple perspectives and potential consequences), and make well-reasoned arguments based on these critiques. Improving communication skills also often fits into these goals, and a number of courses include specific writing or communication-based pedagogy (for example, argumentative writing), even courses based in computer science departments.

Two final types of learning outcomes that exist at a more practical level involve applying rules and creating solutions. For example, some courses have a goal of teaching students to apply philosophical frameworks, legal rules, or codes of ethics to specific fact patterns and situations. Others take the "critique" outcome a step forward towards turning critique into solutions. This is particularly common for courses that seem to be targeted at students who will go on to build technologies—for example, preparing students to "build good technologies," "reduce bad things," or even "[not] destroy the world."

Though there were learning outcomes in our data that did not appear in enough syllabi to become categories for analysis (for example, a few courses included the goal of encouraging students to create their own personal code of ethics), some variation on this recognize/critique/reason cycle was common to most of the courses we analyzed.

5 DISCUSSION AND CONCLUSIONS

Regardless of how severe an ethics crisis within computer science might be [42], this work highlights that academics as a whole are not "asleep at the wheel" [27] when it comes to this problem. University instructors across a range of disciplines have been thinking about what content is most relevant to the topic of technology ethics, and what the goals of that type of ethics instruction should be. Accordingly, discussion of this topic in the CS education research

[1] Based on our analysis in this paper, we have created a list of readings organized into our topic categories, available at www.internetruleslab.com/tech-ethics-syllabi-readings.

community has been slow to start but is on the rise; nearly a quarter of the SIGCSE publications over four decades that mention the word “ethics” have appeared in the past three years.[2] Similar to recent work (e.g., for computing scholarship [38], codes of ethics [39], or ethical principles [15]) that helps parse out what we talk about when we talk about ethics, our findings provide a snapshot of what this looks like in coursework.

Our analysis reveals a great deal of variability across tech ethics courses in terms of the content taught. However, the lack of consistency of content is not surprising considering the lack of standards in this space, and the disciplinary breadth of the syllabi we covered. This is not a bad thing; instead, the variability suggests that there is a lot that computing ethics educators could learn from each other. Having more conversations about this topic in spaces such as SIGCSE can help educators begin to create norms around what computing students should be expected to know about ethics.

However, despite the overall variability, there are still some patterns in these syllabi that reveal specific topics considered critical in this space, such as privacy, algorithms, and inequality/justice. We should begin to consider how these topics intersect with technical content for computer science students and how they might be presented in-situ. Though our analysis included largely standalone ethics classes, there is increasing interest in integrating ethics content into existing technical classes in addition to continuing to teach these ethics-specific classes [16, 23, 30, 36]. Based on the topics that came up consistently in these standalone classes, we can begin to see how these topics might be relevant within technical computer science classes. Some of these are obvious: courses such as algorithms, data science, machine learning, and artificial intelligence are all heavily implicated in tech ethics content. We think that it is critical that these courses, in addition to technical content around issues such as bias and fairness, also touch on the human and social aspects. The topics our analysis uncovered also track directly to courses like design, human-computer interaction, research methods, cybersecurity, and software engineering, as well as courses that might focus on professional and leadership skills for computer science students, like senior project classes.

What might be less obvious is how it might be appropriate to integrate ethics into purely technical programming courses–but this is possible even as early as introductory programming [28]. Our analysis reveals that many topics within tech ethics are high level and conceptual when it comes to the impact of technology on society—e.g., how human decisions are built into code, how technology can reproduce and augment existing social inequalities, how data is created by and directly impacts people, and how choices made at both the level of companies and in small bits of code combine to create large-scale social consequences. This reminder–that code is power, and it should be used responsibility–could be part of every computing course, but is arguably most important at the very beginning of the process of learning to code. This strategy might even be a way to combat an "I'm just an engineer" mindset that ethics is "someone else's job" [12] by emphasizing its role in computing from day one and then continuing this reminder throughout the curriculum.

[2]A search in the ACM Digital Library reveals that out of 6,741 papers published at SIGCSE since 1972, 132 contain the word “ethics” in any field. 30 (or 23%) of these appeared in 2019, 2018, or 2017.

With respect to limitations and future work, we hope that this exploratory, descriptive analysis could be the first step in a broader research agenda for the computing education research community. Because our dataset was not a representative sample of syllabi, we cannot make statements about the overall landscape of tech ethics education, though with purposeful sampling (e.g., by retrieving the syllabi of specific programs [30]), we could explore how ethics currently fits into computing curriculums overall–for example whether specific courses are required. Additionally, as with other syllabi analyses, the actual curriculum described in this study reflects what a large number of instructors consider to be the best ethics course they could offer given resource and time constraints–but we cannot make any claims from this analysis whether any of these strategies are ideal or even adequate [8]. Of course, these self-reported syllabi also might represent those that instructors are particularly proud of or those that belong to social influencers, and therefore allow us to see how this field is likely to be shaped.

There is also space for much more work about the effectiveness of specific types of ethics pedagogy in computer science. Particularly because ethics learning is difficult even for instructors to evaluate in students [25], it is also a challenge to rigorously evaluate ethics teaching effectiveness. However, this is a challenge the SIGCSE community is well-suited to tackle.

Finally, in addition to encouraging contribution to this growing research space, we also hope that this work can serve as a call to action that can encourage and assist instructors at all educational levels who are interested in including ethics as part of their class, as well as computing programs with a goal towards increasing the reach of ethics across a curriculum. The students that we teach to code today will be the ones working at all levels in the tech companies that might—or hopefully might not—be involved in the ethics scandals of tomorrow.

ACKNOWLEDGMENTS

This work was partially funded through the Responsible Computer Science Challenge from Mozilla, Omidyar Network, Schmidt Futures, and Craig Newmark Philanthropies. A huge thanks to Jasmine Bilir, Faisal Lalani, and Rami Mouro for their research assistance on this project.

REFERENCES

[1] Brett A. Becker and Thomas Fitzpatrick. 2019. What Do CS1 Syllabi Reveal About Our Expectations of Introductory Programming Students?. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. Minneapolis, MN, 1011–1017. https://doi.org/10.1145/3287324.3287485

[2] Angela R Bielefeldt, Madeline Polmear, Daniel Knight, Christopher Swan, and Nathan Canney. 2017. Ethics Across the Curriculum? Integrating Ethics and Societal Impact Topics into Core Engineering Courses. In *Proceedings of the ASEE American Society for Engineering Education Rocky Mountain Section Conference*. 1–15.

[3] Judith A. Boss. 1994. The effect of community service work on the moral development of college ethics students. *Journal of Moral Education* 23, 2 (1994), 183–198. https://doi.org/10.1080/0305724940230206

[4] Bo Brinkman and Keith W. Miller. 2017. The Code of Ethics Quiz Show. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. Seattle, WA, 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 (2018), 54–64. https://doi.org/10.1145/3154485

[6] Roxanne L. Canosa and Joan M. Lucas. 2008. Mock Trials and Role-playing in Cmputer Ethics Courses. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. 148–152. https://doi.org/10.1145/1352322.1352187

[7] Cory Doctorow. 2017. More than 50 tech ethics courses, with links to syllabi. *Boing Boing* (November 20 2017). https://boingboing.net/2017/11/20/robo-socrates-what-is-virtue.html

[8] James M. DuBois and Jill Burkemper. 2002. Ethics Education in U.S. Medical Schools: A Study of Syllabi. *Academic Medicine* 77, 5 (2002), 432–437.

[9] Editors. 2018. Cambridge Analytica Controversy Must Spur Researchers to Update Data Ethics. *Nature* 555 (2018), 559–560.

[10] Felicia Chong. 2016. The Pedagogy of Usability: An Analysis of Technical Communication Textbooks, Anthologies, and Course Syllabi and Descriptions. *Technical Communication Quarterly* 25, 1 (2016), 12–28. https://doi.org/10.1080/10572252.2016.1113073

[11] Casey Fiesler. 2018. Tech Ethics Curricula: A Collection of Syllabi. *Medium* (July 5, 2018). https://medium.com/@cfiesler/tech-ethics-curricula-a-collection-of-syllabi-3eedfb76be18

[12] Casey Fiesler. 2018. What Our Tech Ethics Crisis Says About the State of Computer Science Education. *How We Get to Next* (December 5, 2018). https://howwegettonext.com/what-our-tech-ethics-crisis-says-about-the-state-of-computer-science-education-a6a5544e1da6

[13] Howard Gardner. 1991. The Tensions Between Education and Development. *Journal of Moral Education* 20, 2 (1991), 113–125. https://doi.org/10.1080/0305724910200201

[14] Frederick H. Gautschi and Thomas M. Jones. 1998. Enhancing the ability of business students to recognize ethical issues: An empirical assessment of the effectiveness of a course in business ethics. *Journal of Business Ethics* 17, 2 (1998), 205–216.

[15] 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. In *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS)*. 2122–2131. https://doi.org/10.24251/hicss.2019.258

[16] Barbara J. Grosz, David Gray Grant, Kate Vredenburgh, Jeff Behrends, Lily Hu, Alison Simmons, and Jim Waldo. 2019. Embedded EthiCS: Integrating Ethics Broadly Across Computer Science Education. *Commun. ACM* 62, 8 (2019), 54–61. arXiv:1808.05686 http://arxiv.org/abs/1808.05686

[17] R. M. Harden. 2002. Learning outcomes and instructional objectives: Is there a difference? *Medical Teacher* 24, 2 (2002), 151–155. https://doi.org/10.1080/0142159022020687

[18] Ali Haskins Lisle, Coleman Merenda, and Joseph Gabbard. 2019. Using affinity diagramming to generate a codebook: a case study on young military veterans and community reintegration. *Qualitative Research* (2019), 1–18. https://doi.org/10.1177/1468794119851324

[19] Deborah Johnson. 1994. Who Should Teach Computer Ethics and Computers & Society? *Computers and Society* 14, 5 (1994), 6–13. https://doi.org/10.1016/0268-4012(94)90078-7

[20] Kristie J Loescher, Regina W Hughes, Frank Cavico, Jim Mirabella, and Pedro F Pellet. 2005. The Impact of an "Ethics Across the Curriculum" Initiative on the Cognitive Moral Development of Business School Undergraduates. *Teaching Ethics* (2005), 31–72.

[21] Kathleen M. MacQueen, Eleanor McLellan, Kelly Kay, and Bobby Milstein. 1998. Codebook Development for Team-Based Qualitative Analysis. *Field Methods* 10, 2 (1998), 31–36. https://doi.org/10.1177/1525822X980100020301

[22] C Dianne Martin. 1997. The Case for Integrating Ethical and Social Impact into the Computer Science Curriculum. *ITCSE Working Group Reports and Supplemental Proceedings* (1997), 114–120.

[23] C. Dianne Martin, Chuck Huff, Donald Gotterbarn, and Keith Miller. 1996. Implementing a Tenth Strand in the CS Curriculum. *Commun. ACM* 39, 12 (1996), 75–84. https://doi.org/10.1145/240483.240499

[24] David Meyer. 2018. Amazon Reportedly Killed an AI Recruitment System Because It Couldn't Stop the Tool from Discriminating Against Women. https://fortune.com/2018/10/10/amazon-ai-recruitment-bias-women-sexist/

[25] Barbara Moskal, Keith Miller, and L. A. Smith King. 2004. Grading Essays in Computer Ethics. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. 101–105. https://doi.org/10.1145/563517.563380

[26] Norman R. Nielsen. 1972. Social Responsibility and Computer Education. *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education* (1972), 90–96. https://doi.org/10.1145/873684.873708

[27] Cathy O'Neil. [n.d.]. The Ivory Tower Can't Keep Ignoring Tech. *The New York Times* (November 14, [n. d.]). https://www.nytimes.com/2017/11/14/opinion/academia-tech-algorithms.html

[28] Evan Peck. 2017. The Ethical Engine: Integrating Ethical Design into Intro Computer Science. *Medium* (July 5, 2017). https://medium.com/bucknell-hci/the-ethical-engine-integrating-ethical-design-into-intro-to-computer-science-4f9874e756af

[29] Deborah L. Rhode. 1992. Ethics by the Pervasive Method. *Journal of Legal Education* 42, 1 (1992), 31–56.

[30] 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 (2019), 1–26. https://doi.org/10.1145/3341164

[31] Jeffrey S. Saltz, Neil I. Dewar, and Robert Heckman. 2018. Key Concepts for a Data Science Ethics Curriculum. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. Baltimore, MD, 952–957. https://doi.org/10.1145/3159450.3159483

[32] Scott Shane and Daisuke Wakabayashi. 2018. 'The Business of War': Google Employees Protest Work for the Pentagon. *The New York Times* (April 4, 2018). https://www.nytimes.com/2018/04/04/technology/google-letter-ceo-pentagon-project.html

[33] Earl E. Shelp, Michael L. Russell, and Nellie P. Grose. 1981. Students' Attitudes to Ethics in the Medical School Curriculum. *Journal of Medical Ethics* 7, 2 (1981), 70–73. https://doi.org/10.1136/jme.7.2.70

[34] Katie Shilton. 2018. Values and Ethics in Human-Computer Interaction. *Foundations and Trends® in Human–Computer Interaction* 12, 2 (2018), 107–171. https://doi.org/10.1561/1100000073

[35] Natasha Singer. 2018. Tech's ethical 'dark side': Harvard, Stanford, and others want to address it. *The New York Times* (2018). https://www.nytimes.com/2018/02/12/business/computer-science-ethics-courses.html

[36] 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 ACM SIGCSE Technical Symposium on Computer Science Education*. Baltimore, MD, 940–945. https://doi.org/10.1145/3159450.3159573

[37] Michael Skirpan, Jacqueline Cameron, and Tom Yeh. 2018. Quantified Self: An Interdisciplinary Immersive Theater Project Supporting a Collaborative Learning Environment for CS Ethics. *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education* (2018), 946–951. https://doi.org/10.1145/3159450.3159574

[38] Bernd Carsten Stahl, Job Timmermans, and Brent Daniel Mittelstadt. 2016. The Ethics of Computing: A Survey of the Computing-Oriented Literature. *Comput. Surveys* 48, 4 (2016), 1–38. https://doi.org/10.1145/2871196

[39] Luke Stark and Anna Lauren Hoffmann. 2019. Data Is the New What? Popular Metaphors & Professional Ethics in Emerging Data Culture. *Journal of Cultural Analytics* 127, 127 (2019), 1–22. https://doi.org/10.22148/16.036

[40] Manas Tungare, Xiaoyan Yu, William Cameron, Guofang Teng, A P Manuel, Lillian Cassel, Weiguo Fan, and Edward A Fox. 2007. Towards a Syllabus Repository for Computer Science Courses Categories and Subject Descriptors. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. Covington, KY, 55–59.

[41] Marty J Wolf, Don Gotterbarn, and Michael S Kirkpatrick. 2019. ACM Code of Ethics: Looking Back and Forging Ahead. In *Proceedings of the ACM SIGCSE Technical Symposium on Computer Science Education*. Minneapolis, MN, 801–802.

[42] Yonatan Zunger. 2018. Computer science faces an ethics crisis; the Cambridge Analytica scandal proves it. *The Boston Globe* (March 22, 2018). https://www.bostonglobe.com/ideas/2018/03/22/computer-science-faces-ethics-crisis-the-cambridge-analytica-scandal-proves/IzaXxl2BsYBtwM4nxezgcP/story.html