

**Algorithm Bias:
A Survey of Computer Science Student Perceptions**

Abstract

In the United States, Google performs over 3.9 million searches per minute. Monthly desktop searches can exceed over 10.7 billion and mobile searches are predicted to grow steadily. Concurrently, recent discourse has raised questions about bias in search engines and big data algorithms. As the information universe becomes increasingly dominated by algorithms, computer scientists and engineers have ethical obligations to create systems that do no harm. In this paper, the authors discuss a survey that was conducted of computer science and computer engineering students' perceptions of algorithm bias. The aim of the survey was to gather preliminary data on how students perceive bias within machine learning and search algorithms. Over 700 computer science and computer engineering students from three different institutions participated in the survey from Fall 2018 to Spring 2019. Based on survey results, Google was overwhelmingly the preferred search engine. The participants also predicted that artificial intelligence algorithms will improve over time. The majority of respondents believe that private companies, not government organizations, need to regulate their own artificial intelligence algorithms. On average, computer science and computer engineering students acknowledge that algorithm bias could occur when people create algorithms. The results suggest that students are familiar with search engines and in general agreement on how algorithm bias should be addressed in the future.

The survey results will be used to consider whether an information literacy component focused on algorithm bias would be beneficial to offer to students in the computational sciences and if so, how best to design the instruction. The study describes students' prior knowledge for educators seeking to increase awareness of algorithm bias. Our hypothesis is that computer science student exposure to the concept of algorithm bias via instruction would create positive changes in the technology workforce as students with training in algorithm bias mitigation bring their knowledge to the sector. A commitment to understanding and reducing algorithm bias in the tech industry would create spaces where communities can optimize their search for information and expect fair treatment from automated systems.

Introduction

Recent discourse in information ethics has raised questions about bias in search algorithms and machine learning. Algorithms are sets of instructions within computer programs that direct how these programs read, collect, process, and analyze data. Algorithms have become part of the architecture of much of the internet and are also the basis of artificial intelligence (AI). We use the term algorithm bias to refer to computer systems that “systematically and unfairly discriminate against certain individuals or groups of individuals in favor of others.”¹ Several

articles and books, such as Safiya Noble's *Algorithms of Oppression* (2018)² have discussed the phenomenon.

There are many reasons why an algorithm may be considered "biased." Incomplete or faulty data is one reason. For example, in a study published in *Nature Communication*³ researchers confirmed for the first time that two of the top genomic databases, which are in wide use today by clinical geneticists, reflect a measurable bias toward genetic data based on European ancestry over that of African ancestry. This deficit in African ancestry genomic data was identified during an 18-month long study conducted via the Consortium on Asthma among African-Ancestry Populations in the Americas (CAAPA). When compared with current clinical genomic databases, researchers found a clearer preference in those databases for European genetic variants over non-European variants. This gap in genomic data can impact how disease risks are analyzed for people of African descent vs. those of European descent. Awareness that data may be skewed toward a particular group or exclude entire groups is necessary for correcting problems with datasets; machine learning alone won't be able to rectify an incomplete dataset. Humans have to be aware of racial, gender, or other discrepancies, first, in order to be able to identify data as incomplete or in need of correction.

Another reason for algorithm bias is the possibility of bias inserted by humans. An instance where human bias may be at play is that of Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), a case management tool widely used in the United States to guide sentencing by predicting the likelihood of a criminal reoffending. In May 2016, ProPublica⁴ reported that the COMPAS system predicts that black defendants pose a higher risk of recidivism than they actually do, and the reverse was predicted for white defendants. Equivant, the company that developed the software, disputes this discrepancy. It is hard to pinpoint where the bias, if any, might come from, because the algorithm is proprietary.

Those who have studied the phenomenon of algorithm bias posit that the workforce that creates algorithms that shape our information infrastructure remains predominantly male and white; the hypothesis that this skew in gender and race affects algorithm design has been backed by evidence from computer and data scientists studying the phenomenon.⁵

As computer science educators and librarians who work regularly with computer science and engineering students, we strongly recommend harm mitigation via societal self-reflection about artificial intelligence. More specifically, we suggest that our students who will become future developers of algorithms and artificial intelligence need a comprehensive introduction to ethical considerations in algorithm development. Our study examines how computer science students perceive algorithm bias in order to build on their prior knowledge for instructional

considerations. Effective instruction involves understanding what students know and building on their existing perceptions and experiences.⁶

The existing surveys of algorithm bias focus on general populations. Worldwide 46% of people believe search engine algorithms are unbiased and 38% of people in the United States agree the algorithms to be unbiased.⁷ When asked about the causes of bias and “unbiasedness of algorithms,” people mentioned the absence of human input. In other words, humans add both objectivity and bias. Along the same lines, a Pew Research Center⁸ survey points out another complexity: 58% of Americans believe that AI algorithms will always reflect human bias and 40% of survey participants agree that AI can also become unbiased if properly designed by humans.

Objectives and Methodology

Our survey of prevailing attitudes about algorithm bias focused on undergraduate and graduate computer science and computer engineering students. The aim of the survey was to understand how students see search engines and bias within machine learning algorithms. Although surveys on algorithm bias have been conducted in the past, the computer science student population has not specifically been studied. The authors wanted to learn more about the student population that will analyze, code, and design future algorithms. What do students know about search engine and machine learning algorithms?

This paper focuses the following research questions:

1. What are computer science students’ experiences with search engines?
2. What are their beliefs about search engines and algorithms?

From December 2018 to March 2019, we surveyed computer science students via a 10-minute online survey that included questions on search engines, technological optimism, and ethics education. Recruitment methods for the survey included flyers and in-person and email communications with computer science faculty and staff who coordinate undergraduate, graduate, and online programs. We made announcements about the survey during library instruction sessions in computer science and other programs. In the end, we found that newsletters and online course management systems also increased participation.

The cross-institutional study between three institutions have distinctly different student populations. The University of Southern California (USC) has a diverse student body in terms of race and gender. Women comprised 44% of a recent incoming class.⁹ California State University, Los Angeles (Cal State LA) is a Hispanic-Serving Institution (HSI) where first

generation students are almost 58% of the student population.¹⁰ Boise State University is more racially homogeneous but has seen an increasing number of non-traditional aged students in attendance. The majority (78%) of responses are from with USC remaining responses evenly divided between Boise State and Cal State LA.

The following table shows the participation rates by university.

University	No. of responses	% of total
Boise State University	81	11%
Cal State LA	77	11%
USC	566	78%
Total	724	100%

Table 1. Institution responses by percentage.

The majority of participants were graduate students (63%) and remaining (27%) were undergraduate computer science students. The margin for sampling error for the responses considered is ± 3.2 percent at a 95 percent confidence interval. Margins of error for subgroups within the survey sample will be higher. For example, the margin of error for female students would be ± 6.34 percent. Opinions of the respondents regarding the questions on search engine results and algorithm bias were recorded in the form of a 7-point Likert scale ranging from “Strongly disagree” to “Strongly agree.” The ordinary category variable has been converted to numerical ranging from -3 to +3 and mean and standard deviation is measured for the overall opinion and its distribution and also for the subgroup opinion and its distribution.

Students answered a variety of survey questions about their experience using search engines, whether they had positive or negative experiences while using search engines, and their beliefs about AI. They were not asked directly about algorithm bias or asked to define the term. We

designed the survey with the primary goal of identifying students' experience, prior knowledge, and perceptions about search engines and AI.

A graduate student assistant, Adarsh Gopalakrishman (USC), majoring in data science was hired to help us with survey publicity and later clean and compile data after the survey period closed. The student assistant helped prepare the survey data (i.e., remove incomplete and inaccurate records) and organize the data for analysis. It should be mentioned that our study is not intended to be a complete formal quantitative investigation; it serves as a consideration of students' prior knowledge for effective instruction. Validation of the results with larger studies and classroom teaching may be required.

Results

Research Question 1: What are the computer student experiences with search engines?

In the survey, students were asked about their search engine activities and experiences. The computer science students search engine preferences and weekly usage were clear. Google was the dominant search engine by a 97% margin. Additionally, 88% of respondents indicated they use computer search engines daily. The student responses suggest Google searches are a daily habit.

Interestingly, when considering the quality of search engine results, student confidence levels decreased as the evaluation criteria or standards increased from everyday use to trustworthiness. We see 94% agreement gradually falling to 44% agreement as the standards for search engine results rise.

Student agreement on search engine results by percentage	
Search engine results are good enough for every day questions.	94%
Search engine results are accurate.	80%
Search engine results are complete.	58%
Search engine results are trustworthy.	44%

Table 2. Student perceptions on search engine results.

Most people surveyed do not necessarily trust search engines; 56% are neutral or negative about trustworthiness.

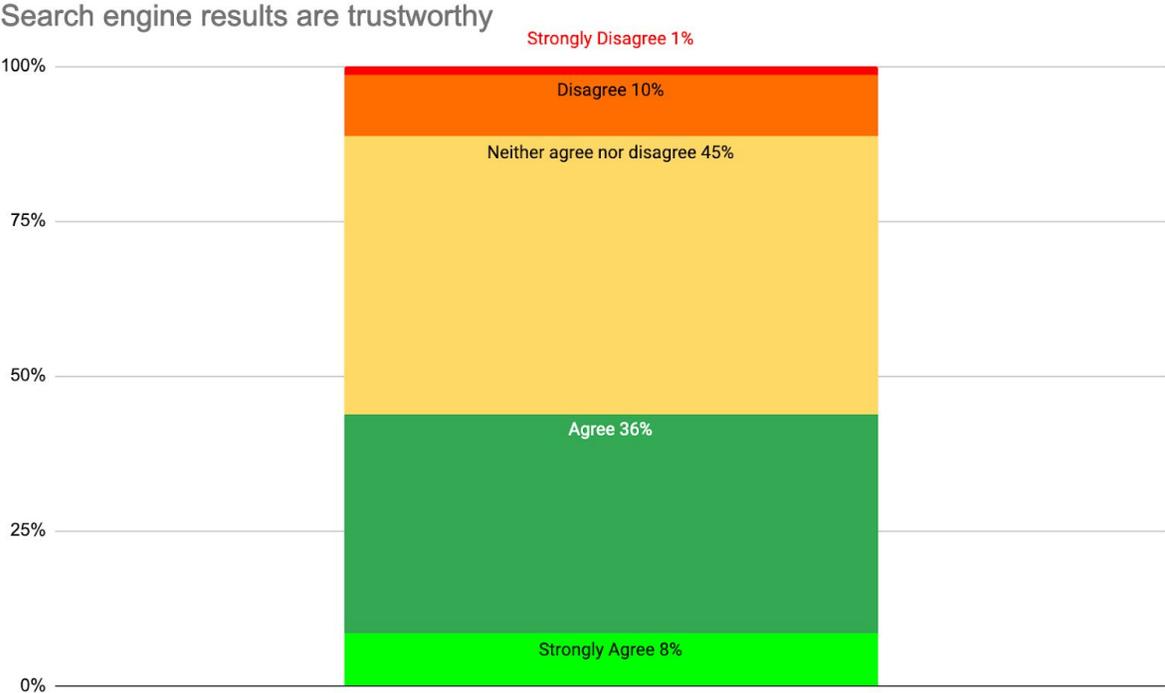


Figure 1. Students’ perceptions of search engine trustworthiness.

In conclusion, most students frequently use Google as their primary search engine. They seem satisfied with Google searching for accuracy and results for everyday questions. A limited number of students question or remain neutral with respect to the trustworthiness of results.

Search engine results are

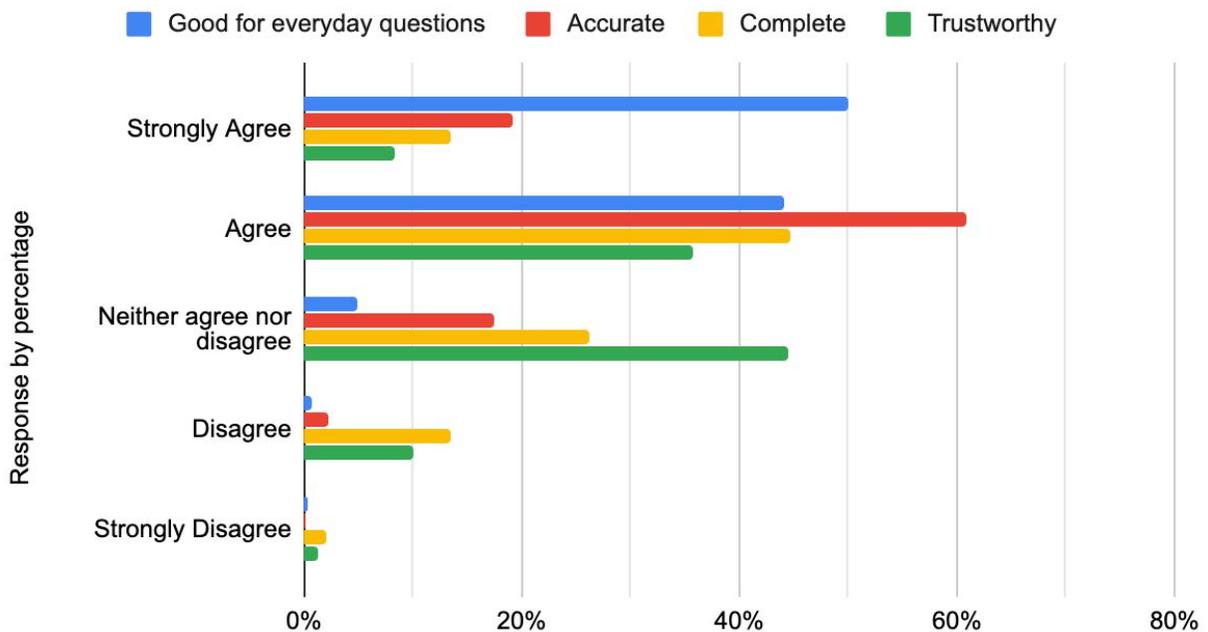


Figure 2. Students' perceptions of search engine results.

Research Question 2: What are student beliefs about search engines and algorithms

As we consider student beliefs about search engines and algorithms broadly, we find additional positive impressions. Computer science students seem optimistic as an overwhelming 92% of student participants believe that artificial intelligence algorithms will improve over time. Along the same lines, the majority of students (83%) expect private organizations to be responsible for ensuring that they serve the general public.

Conversely, opinions about government regulation of search results are more varied: 52% of students approve of government regulations, 20% are neutral, and 28% disapprove. Although most people accept the need for government regulations, a considerable minority disagree about the government's regulatory role.

Oversight

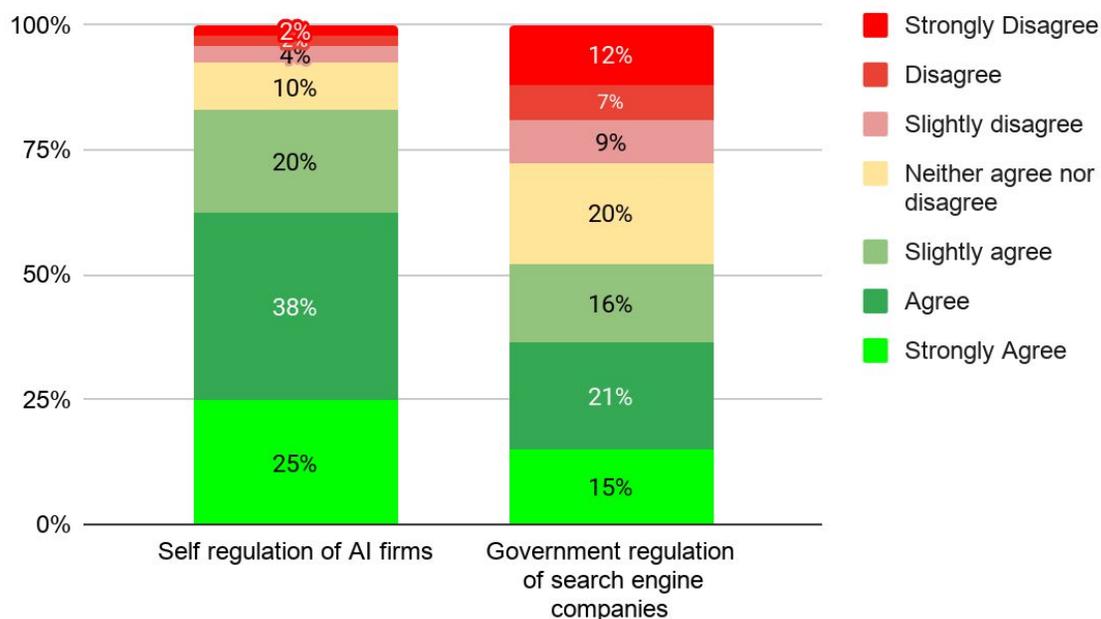


Figure 3. Students’ perceptions of AI companies.

Where else did students agree on algorithm bias? On average, computer science and computer engineering students (83%) acknowledged that algorithm bias could occur when they responded to the statement “Engineers can create biased algorithms.”

From the survey responses, several key observations can be made with respect to computer science student perceptions. Although students surveyed share some concerns, their largely positive perceptions support an optimistic view of technology being able to tackle algorithm bias. Students report frequently searching with Google and finding accurate results good enough for common questions. Computer science students also shared their substantive optimism about the future of AI code and confidence that technology companies can self-regulate. A belief that technology can be harnessed to help make the future better was consistently presented. It might be worth following up, if a future survey were to be conducted, on why students feel confident AI and machine learning algorithms will improve.

Simultaneously, 83% of our respondents agreed that “engineers can create biased algorithms” and that some search engine results were not necessarily entirely complete and trustworthy. Their opinions about the role of government regulation for fair search engine results could be interpreted as divided given that 52% agree, 28% disagree, and 20% remain neutral. Their

reservations concerning engineers who create biased algorithms and incomplete or questionable or unreliable search results may suggest more nuanced opinions about algorithmic bias on some topics in specific areas.

Discussion

Computer science students present a decidedly optimistic view of the future of AI code; a belief that technology can be harnessed to help make technology better was consistently presented. A significant majority also realize that engineers can create biased algorithms. These two responses in combination would seem to indicate the possibility that those being trained to develop AI code recognize the possibility of bias but hold an unwavering optimism that could reasonably lead them to overlook the potential of bias in their professional work. The optimism is also expressed that for-profit organizations or companies that create AI and automated tools will voluntarily self-regulate in society's best interest.

Although the students we surveyed agree companies should self-regulate, more commonly, state, local and federal government and other parties will investigate discrimination of protected classes for "accuracy, fairness, bias, discrimination, privacy, and security." Moreover, scrutiny of algorithmic discrimination at a federal level could be strengthened should the Algorithmic Accountability Act of 2019¹² become a law. A recent example that received media scrutiny is the possible gender discrimination in the Apple Card's differing credit limits for men and women, even when credit histories and incomes were identical. Including such contemporary case studies as part of a computer science ethics curriculum could develop students' critical thinking skills to explore their prior knowledge and enhance learning about the role of regulation and oversight in AI development.

This study's aim is to understand students' perceptions on search engine algorithm bias in order to articulate future steps for effective education. Outlining students' prior knowledge aids in developing effective instruction modules⁶. Based on the overwhelming technology optimism and a few uncertainties/contradictions, such as doubt that search results are completely trustworthy and recognition that engineers can be biased, we speculate that students in the sample may share some of the sentiments of the general public, as articulated by various reports,^{7,8} finding the emerging technology complex and contextual.

We are aware that students may have differing definitions of algorithm bias. Understanding of the phenomenon of algorithm bias is rapidly evolving and thus definitions and scope may change. Presumably, the survey participants may have different interpretations of algorithm bias and perhaps, even difficulties articulating a definition. Although there are standard definitions in the literature,¹ students may not be familiar with them yet.

Moreover, the survey focused on the students' perceptions rather than definitional awareness. Their impressions of search engines, AI, and other relevant algorithm bias topics serve as starting points for discussions and development of a more formal instructional module on algorithm bias.

A more open-ended discussion early on in an instructional module can enable educators and students to develop a shared understanding of algorithm bias and its impact on future computer science professionals. Although discussing what is and is not algorithm bias for students vs. computer science experts (professors and researchers) is foundational, we also suggest that, in addition to creating a shared working definition, engaging students with ethical considerations (e.g., search engine bias, regulations, community values and practices, and costs) in computer science assignments and challenges could connect to students' prior knowledge and contextualize the importance of professional ethics. The interactions have the potential to engage critical thinking, highlight multiple perspectives, all to improve instruction and student learning in ethics education. In other words, given that students' mental models are different from expert instructors, the survey findings may help instructors become more effective or aware as they incorporate micro-lessons, modules, or problem-based learning in their curriculum.

What else might engineering and computer science faculty and librarians do? Asking questions about student experiences could help educators gauge how to frame a lesson plan and clarify learning objectives. From there, relevant learning objectives such as recognizing definitions and examples of algorithm bias and its harms or additional foundational concepts could emerge. An alternative approach could include having students write their own definitions of algorithm bias and its impact as a method of engaging their prior knowledge and mindsets.

Including recent articles or readings in activities fosters informed discussions of the harms and solutions, possibly highlighting that knowledge creation can be contextual, non-linear, and iterative. In other words, information, research, and theories can be contradictory, even complex and thus, needs to be debated before becoming widely accepted. Additionally, educators could ask students to brainstorm solutions and interventions based in different arenas or perspectives or identify relevant frameworks such as ethical, philosophical, educational, social, corporate, government, professional organizations, regulatory, and so on.

This study is an initial effort to understand the mindset of computer science students with regard to search engines and algorithm bias. The findings serve as a starting point for creating a data-driven introductory module on algorithm bias. We are currently testing and refining the instruction module across three institutions. By testing across three institutions, we collected comprehensive feedback from diverse student populations to further develop our understanding of computer science students' mindsets. Because librarians value how information is produced,

evaluated, and accessed, they can play a key role in partnering with engineering educators to help students participate in conversations about algorithm bias.

Reflections

We did not want to skew the survey by asking students about algorithm bias directly. There were no questions asked in the survey that used the term algorithm bias. However, students were asked questions about their experiences with search engine use and their opinions about related areas including ethics training. Most participants agreed that they were well trained about the ethical impacts with respect to designing computer algorithms. The differences between graduate and undergraduate students were negligible. The difference between upper- and lower-division undergraduate students who disagreed with the effectiveness of ethical impacts was surprising (13%) and could be an area for future research.

Computer science students are well trained about the ethical impact of their technology design choices for computer algorithms.			
	Disagree	Neutral	Agree
Graduates students (n=454)	25%	11%	64%
Undergraduates students (n=266)	25%	14%	62%
Undergraduate students in 1st and 2nd years (n=146)	19%	15%	66%
Undergraduate students in 3rd and 4th years (n=120)	32%	12%	57%

Table 3. Students’ perceptions about ethics training

Some educators may argue that the existing computer science ethics education does not impact student views on algorithm bias and their professional knowledge. Indeed, taking one or possibly two engineering ethics courses as part of a degree program may not have a lasting impact on students embarking on professional careers. However, alternatives to semester-long ethics courses may be more effective and impactful. In a systematic literature review on engineering ethics education, Hess¹³ (2018) suggests integrating micro-insertions of ethical decision-making across the curriculum. For example, the micro-insertion of asking students to choose between two refrigerants by weighing not only environmental impacts but also adding cost considerations¹⁴ could increase awareness that “ethical considerations are present in every event.”¹³ Hess also stresses that in addition to inserting practical engineering ethics considerations into existing assignments, student participation in problem-based, community

engagement projects such as Engineers Without Borders¹⁵ reinforces “a holistic understanding” of engineering ethics.¹³ As students appreciate how ethics grounds their practice, they unknowingly practice the ethics of care.^{16,17} They care for stakeholders and consider multiple perspectives. The promising engineering instructional practices and community-based learning can be extended to computer science curricula.

Currently, we are in a developmental phase for the instructional component on algorithm bias based on the survey analysis. One limitation we face is that our survey data covers a relatively small set of questions. In future research, we would like to gather more qualitative data as well as expand the scope of questions. We hope to develop evidenced-based instructional activities to help students become more aware of ethical considerations when designing automated decision making systems.

Bibliography:

- [1] Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems (TOIS)*, 14(3), 330–347. <https://doi.org/10.1145/230538.230561>
- [2] Noble, S. (2018). Algorithms of oppression: How search engines reinforce racism.
- [3] Study Reveals Major Racial Bias in Leading Genomics Databases. (2016, October 11). Laboratory Equipment, p. n/a. Retrieved from <http://search.proquest.com/docview/1830956349>
- [4] Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2018). Machine Bias. *Nieman Reports*, 72(3/4). Retrieved from <http://search.proquest.com/docview/2136026806/>
- [5] West, S.M., Whittaker, M. & Crawford, K. (2019). Discriminating Systems: Gender, Race and Power in AI. AI Now Institute. Retrieved from <https://ainowinstitute.org/discriminatingystems.htm>
- [6] Bransford, J. D., Brown, A. L., & Cocking, R. R. (2000). *How people learn* (Vol. 11). Washington, DC: National academy press. <https://www.nap.edu/read/10129/chapter/8#118>
- [7] CIGI-Ipsos 6. (2019). *2019 CIGI-Ipsos Global Survey—Part 3 Social Media, Fake News & Algorithms*. Retrieved January 14, 2020, from <https://www.cigionline.org/sites/default/files/documents/2019%20CIGI-Ipsos%20Global%20Survey%20-%20Part%203%20Social%20Media%2C%20Fake%20News%20%26%20Algorithms.pdf>
- [8] Smith, A. (2018). *Public Attitudes Toward Computer Algorithms*. Retrieved from <https://www.pewresearch.org/internet/2018/11/16/public-attitudes-toward-computer-algorithms/>
- [9] Ballon, M. (2017, October 9) Incoming undergrad class at USC Viterbi is 44 percent female—a school record. USC News. Retrieved from <https://news.usc.edu/128553/the-number-of-female-engineers-at-usc-viterbi-grows/>
- [10] Institutional Effectiveness. (2020.). Retrieved from March 10, 2020, <http://www.calstatela.edu/InstitutionalEffectiveness/student-enrollment>
- [11] Barkho, G. (2019, November 16). How Goldman Sachs Can Regain User Trust After Apple Card Discrimination. *Observer*. <https://observer.com/2019/11/goldman-sachs-bias-detection-apple-card/>
- [12] Clarke, Y. (2019) Algorithmic Accountability Act of 2019. Retrieved from <https://www.congress.gov/bill/116th-congress/house-bill/2231/text>

- [13] Hess, J. L., & Fore, G. (2018). A Systematic Literature Review of US Engineering Ethic Interventions. *Science and Engineering Ethics*, 24(2), 551–583. <https://doi.org/10.1007/s11948-017-9910-6>
- [14] Davis, M. (2006). Integrating ethics into technical courses: Micro-insertion. *Science and Engineering Ethics*, 12(4), 717–730.
- [15] Wittig, A. (2013). Implementing problem based learning through Engineers Without Borders student projects. *Advances in Engineering Education*, 3(4), 1–20.
- [16] Hedayati Mehdiabadi, A. (2018). The Ethical Judgement Processes of Students in Computing: Implications for Professional Development Paper presented at 2018 ASEE Annual Conference & Exposition, Salt Lake City, Utah. <https://peer.asee.org/31099>
- [17] Tronto, J. (2005). An ethic of care. In A.E. Cudd, & R.O. Andreasen (Eds.) *Feminist theory: a 1414 philosophical anthology*, (251-263), Oxford: Blackwell Publishing.