How do Undergraduate Students Reason about Ethical and Algorithmic Decision-Making?

Ashish Aggarwal
ashishjuit@ufl.edu
Department of Engineering Education
University of Florida
Gainesville, FL, USA

Saurabh Ranjan
saurabh.ranjan@ufl.edu
Department of Psychology
University of Florida
Gainesville, FL, USA

ABSTRACT

As the effectiveness of algorithms to make decisions improves and as the use of algorithms in domains, which can have a significant impact in determining one's life prospects increases, it is important to understand undergraduate students' perceptions of algorithmic decision making and reasoning behind that perception. We conducted a study to understand engineering students' perception about algorithmic decision making in two different scenarios using a trolley problem at the end of an introductory programming course. The motivation to conduct this study was to gain insights on how they reason about the ethical use of algorithms. Data of eighty-two undergraduate engineering students was analyzed to not only understand their decisions in two different contexts but also their qualitative reasoning behind their decisions. This paper presents a thematic analysis of these decisions and how they differed in the two contexts. Further, classification of their reasoning into different known philosophical frameworks is discussed, which helps in understanding the major underpinnings of these decisions. We believe that the results of this study can help educators understand how students reason about algorithms which may influence how "ethics" as a topic is integrated in computer science courses, especially in introductory programming courses.

CCS CONCEPTS

• **Social and professional topics → Computer science education**; **CS1**; *Codes of ethics.*

KEYWORDS

CS1, algorithmic decision making, ethics, computational reasoning, trolley problem

ACM Reference Format:
Ashish Aggarwal and Saurabh Ranjan. 2022. How do Undergraduate Students Reason about Ethical and Algorithmic Decision-Making?. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2022), March 3–5, 2022, Providence, RI, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3478431.3499412

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

ACM ISBN 978-1-4503-9070-5/22/03…$15.00
https://doi.org/10.1145/3478431.3499412

1 INTRODUCTION AND MOTIVATION

As the use of algorithms increases in our everyday life, and as the effect of these algorithms in impacting important individual outcomes, it has become important to discuss and cultivate a sense of ethics in our undergraduate computing curriculum [13, 14]. There is an increasing amount of research and resource creation on integrating ethics in computing courses, especially in introductory programming courses [16]. While there are multiple pedagogical approaches for integrating ethics in CS curriculum, like an explicit course based approach or covering it across the multiple courses [5], most of the curriculum is informed by a theoretical moral philosophical underpinnings guided by experts. As the nature of ethics and how it is perceived with respect to changing technological developments becomes more complex [3], it would be valuable to understand students' default reasoning patterns about ethical and algorithmic decision making processes. We believe such research can inform and guide curriculum and educational perspective to make ethical-CS learning more engaging and participatory. To address and advance research on these lines, the paper presents the results of a study which uses a grounded theory approach to explore undergraduate students' decisions and reasoning in an ethically challenging situation. The results of this study will help inform CS-Ethics curriculum design by taking into account students' default ethical reasoning patterns.

2 BACKGROUND

Multiple studies have looked at students perceptions about ethical issues at the intersection of engineering, philosophy, politics, law, sociology etc., [7]. To understand ethical issues in computing, Salehnia et al. studied students' perceptions using thirty questions based on their ethical behavior in regard to oblige, opportunities, intent and professional responsibilities' across CS and non-CS students [15]. Similarly, Freyne et al. studied engineering students' perception of engineering cases involving advances in transportation, agriculture, space, weapons of mass destruction, video games etc., [7]. They also did a comparative study to understand engineering and non-engineering students' perceptions of ten ethical issues. They presented similarities and differences in students' perceptions, highlighting students' viewpoints on multiple cases [8]. Ethics have various dimensions, and while these studies focused on technological advancements and its effect on society at large, none of them explicitly studied students' perspectives on algorithmic decision making. Also, while multiple studies have suggested ways to improve ethics curriculum [11], there is a lack of research on integrating undergraduate students' ethical viewpoints in informing curriculum design.

There is now a growing body of literature that is studying the approaches to integrate ethics in the CS curriculum [5, 10]. Doore et al. and Skirpan et al. have suggested multiple activities and case studies to incorporate ethics in CS courses [6, 16]. In studying the ethical issues associated with privacy in robots, Lutz et al. suggested involving multiple stakeholders to analyze ethical implications in reviewing big data analysis and predictions [12].

As a category of problems that help understand and contemplate ethical issues in autonomous systems, trolley problems have been widely used by researchers and AI-educators [9]. A simple version of a trolley problem can be described by an ethical-dilemma in which a self-driving car's breaks fail, and it has to decide what to do (Fig: 1). The dilemma firstly comes from whether the algorithm should even choose to do something in such a case or not. And if yes, then what should it prioritize? Saving passengers in the car or others who may be nearby. This increasingly becomes complex as there is more context added to it. Its various versions have been used by philosophers and scientists to understand the ethical issues related to automated decision-making. In a study conducted to understand the dilemma of autonomous vehicles and what they should do, Bonnefon et al. *"...found that even though participants approve of autonomous vehicles that might sacrifice passengers to save others, respondents would prefer not to ride in such vehicles"* [4]. Most recently, Awad et al. worked on understanding the global variations in ethics using the different factors in the versions of the trolley problem. In an online experiment where they collected data from all over the world, they found that firstly, moral choices are not universal, and secondly, contextual cultures, socio-economic conditions influence a person's decision [1].

The trolley problems have also been criticized primarily for being disconnected from reality [2]. With all the limitations and shortcomings of the trolley problem, it is one of the most commonly discussed examples, guiding the debate on the ethical issues. It is one of the more practical examples, which can be used to introduce the complexity of the decision-making processes in an undergraduate level class or any general discussion on AI-Ethics [9]. Our goal in this paper is to understand how undergraduate students reason about ethical and algorithmic decision-making. Thus, to answer our research questions, we use the trolley problem examples to better understand students' decisions and reasoning about ethical decision-making in two different conditions. Note that here we are neither assessing the use of trolley problems in ethics-curriculum nor studying its effectiveness. We are simply using it to understand students' decision-making processes to explore their reasoning and moral underpinnings. The study results can inform and guide the creation of an ethics curriculum where students' reasoning patterns are taken into account.

3 RESEARCH QUESTIONS

As discussed above, we are interested in exploring undergraduate students' reasoning on ethical and algorithmic decision-making. To do this, we use a version of the trolley problem where there is a car whose breaks have failed. There is a truck ahead which is slowing down, there is a motorcyclist without helmet on the left and another motorcyclist with helmet on the right (Fig: 1). The dilemma is to decide what the car should do. To understand students decisions, we create two conditions. In condition-1, students are asked to decide and explain what they will do if they are driving the car and they find themselves in this situation. In condition-2 students' assume the role of a programmer who writes an algorithm for a self-driving car for this particular situation and discusses their decisions. We are interested in knowing how their choices and reasoning in condition 1, where they are the driver, contrasts with their choices and reasoning when tasked with writing an algorithm which will be integrated at scale. Specifically, the research questions are:

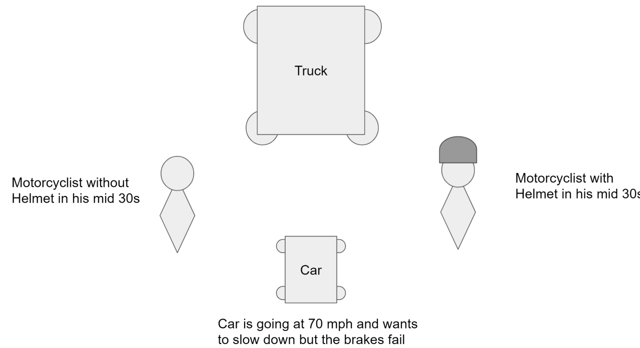

Figure 1: A visual description of the scenario for both the conditions used in the survey

(1) How do undergraduate engineering students reason about ethically challenged circumstances:
  (a) Condition 1: When the student is driving the vehicle?
  (b) Condition 2: When the student is designing algorithm for operation without human oversight?
(2) How does their reasoning change from condition 1 to condition 2?

4 METHODS

In an introductory programming course for engineering students in Summer-2018, a survey was used, and responses from 82 students were collected. The survey was an extra credit assignment given at the end of the course, and it had multiple questions out of which two main questions related to this analysis were presented. In the scenario, **a car is going at a speed of 70 mph. A truck is in front of it, headed in the same direction but slowing down quickly (Fig: 1). To prevent a collision, the car needs to slow down, but its brakes fail**. On the left-hand sight there is a motorcyclist in his thirties **without a helmet (left side)** who is driving the motorcycle at 70 mph. On the right-hand side also there is a motorcyclist **with a helmet (right side)** in his thirties who is driving the bike at 70 mph. In the first condition (Condition-1: SELF), the students are asked, assuming they are driving this car, what will they do in this scenario and why? In the second condition (Condition-2: ALGORITHM), students are asked to propose an algorithm to govern autonomous vehicle decision-making in the same scenario, assuming that they are writing an algorithm for a self-driving car company as an intern.

This study was approved by our University's Institutional Review Board under the Exempt category. Out of 82 students, 62% *(n = 51)* of the students were male, and 38% *(n = 31)* of the students were female. Overall, 45% of the students had mechanical and aerospace engineering as their major, 20% had biomedical engineering and the remaining 35% were from other majors including other engineering majors. This was the first programming class for approximately 60% of the students.

To obtain students' generic decisions and reasons, the survey had essay type free-response text box for both the conditions. We used a grounded theory approach to identify the themes and then color code students' reasoning to identify a dominant reason for every decision. Two researchers, including the first author, did two separate qualitative codings, and the results were compared for inter-rater reliability. Any difference was resolved by mutual discussion. This methodology was preferred over giving some options on the survey because we wanted to qualitatively mine students' reasoning rather than prompting them to specific reasons they may or may not think about while submitting their responses. In this way, we tried to reduce the influence on their reasoning.

Based on the various students' reasons, we broadly categorized the reasons into eight themes described in Table:1. Some of the reasons are not mutually exclusive, and sometimes students' responses indicated multiple reasons. However, for analysis purposes, we selected a dominant reason based on a students' response after carefully considering its explainability for the decision.

Table 1: Themes of major reasons used in students' reasoning and their meaning

Themes for Reasons	Description
Other	These were the reasons which were not able to be categorized in any other reasons and were outliers or not clear
Whatever causes least causality	It refers to the minimization of overall harm caused in the scenario indicating a utilitarian approach
I'm Responsible	This reason indicated an approach whose decision was based on the fact that the driver of the car was responsible for the situation as they were driving the car whose breaks have failed
Protect Myself	This indicated the prioritization of one's life over any other consideration
Customer-Client	This refers to liability implications associated with a company which has certain legal obligations to its customers
Cars are Safer	Multiple reasons relied on the possibility of car's safety features like airbags or physical properties which made cars safer even in the event of an intense collision
Algorithms should not decide	This indicated a more principled reason where the very fact that algorithms can decide who to harm and who not to harm, is fundamentally disapproved
All Lives are Equal	This reason was based on not prioritizing any individual's life (especially that of their own) when making the decision
Presence of Helmet	In this reason, the presence (or absence) of a helmet was a critical consideration which influenced the decision

5 RESULTS

On analyzing students' decisions in Condition 1: SELF (Fig: 2), we found that 10% of students (n=8) chose to swerve to the left (towards the motorcyclist without helmet), 30% of students (n=31) decided to swerve to the right(towards the motorcyclist with helmet), 45% students (n=37) decided to go straight, and 7% (n=6) had other decisions like use the parking breaks or swerve in the direction of less traffic, which were unable to be categorized in one of the major decisions. On further analyzing the themes based on their reasoning, it was found that 39% (n=32) students' decisions were influenced by the presence (or absence) of the helmet and 24% of students' decisions relied on safety features in the car while making a decision (Fig: 3).

On analyzing students' decisions in Condition 2: ALGORITHM (Fig: 4), we found that only 4% of students (n=3) chose to swerve to the left (towards the motorcyclist without helmet), 17% of students (n=14) chose to swerve to the right(towards the motorcyclist with helmet), 22% students (n=18) decided to go straight, 38% students (n=31) decided that their algorithm will leave it to the driver to decide and 19% (n=16) had other decisions. On further analyzing the themes it was found that 30% (n=25) of students' decisions were influenced by the reason that algorithms should not take any action in such scenarios. 20% (n=16) students indicated that their decision would cause least casualties. 18% (n=15) students' decisions relied on the motorcyclist wearing (or nor wearing) helmets (Fig: 5).

While these were the aggregate results based on students' responses, we are interested in further exploring subgroups of students responses and their differences (if any) in two conditions. For this purpose, we analyze students' responses in condition 2 by grouping their decisions for condition 1. So the following section discusses three analysis based on students' decisions in Condition 1: SELF. We first discuss the students who chose to swerve left (towards the motorcyclist without helmet) in condition 1 and analyze their choices and changes in reasoning in condition 2. This is followed by the analysis of group of students who chose to swerve right and who indicated to go straight in condition 1. Such analysis will help us explore more nuanced results, which will help answer the research questions.

6 FINDINGS

6.1 Analysis 1: When students chose to swerve 'LEFT' (towards the motorcyclist without helmet) on Condition1: SELF *(8/82)*

There were eight students who chose to swerve left in condition 1. Out of those eight, seven indicated that the reason for swerving left

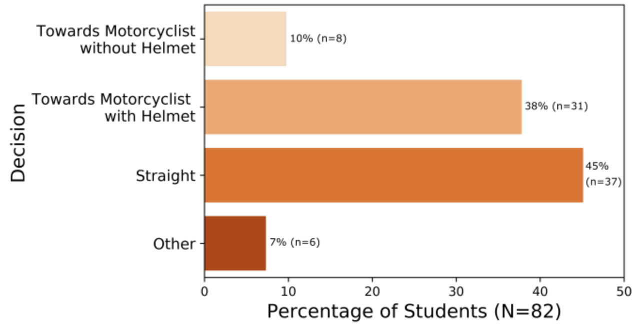

Figure 2: Students' decision on Condition 1: SELF when they are themselves driving the car

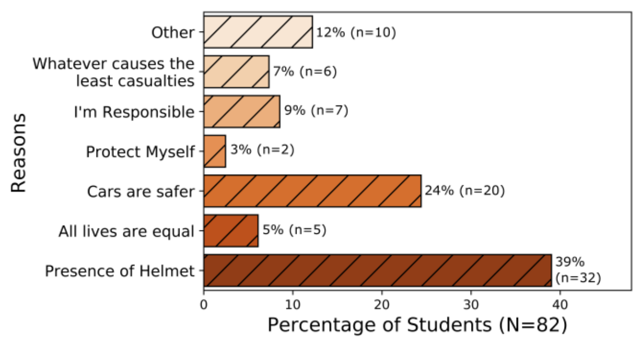

Figure 3: Students' reasons on Condition 1: SELF when they are themselves driving the car

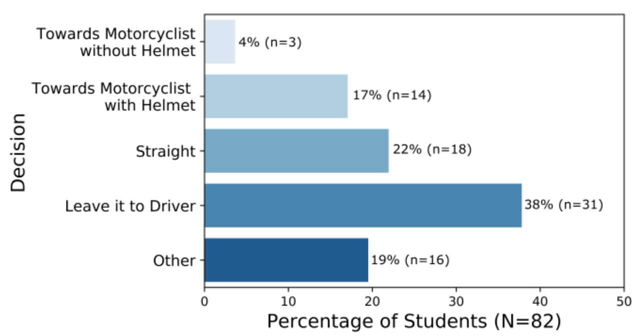

Figure 4: Students' decision on Condition 2: ALGORITHM when they are writing an algorithm

was the motorcyclist not wearing helmet. An example of a student's response is the following:

> *"I would swerve to the left and hit the motorcyclist without the helmet. While the human race depends on natural selection to improve the gene pool, artificial selection can be just as effective. Not wearing a helmet on a road with a speed limit greater than or equal to 70 mph is simply irresponsible and eliminating that negligence is the best case scenario"*

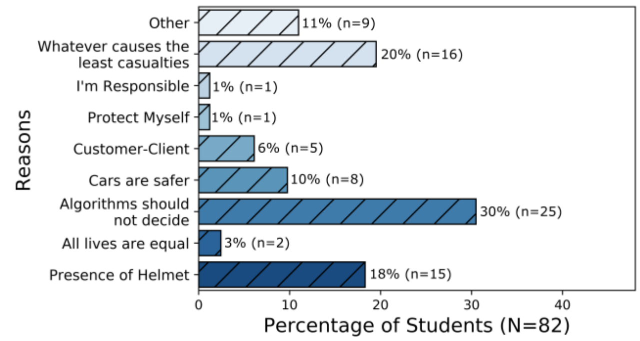

Figure 5: Students' reasons on Condition 2: ALGORITHM when they are writing an algorithm

This suggests that the students chose to go to the left because the motorcyclist on the left did not wear a helmet, implying either he was not responsible enough or did not care for his life. In philosophical framework, this would be an example of a decision based on *retributivist justice* which means that "when an offender breaks the law, justice requires that they suffer in return, and that the response to a crime is proportional to the offense" [19].

However, when the same students were asked to make a decision in the algorithmic context, three out of eight said it should be left to the driver because, primarily, algorithms should not make such a decision. This indicates that approximately 38% (3 out of 8) adjust their decision in algorithmic context and their reasoning is influenced by the fact that algorithms should not make such decisions. The following is a response from a student who says that there are potential legal issues involved with algorithms making such decisions:

> *"I would propose that the car does not take action in this case, leaving everything to the passengers in the car. The passengers should be given a warning, but ultimately the car would not take action, largely due to a host of potential legal issues that the company could face following the accident (damages, charges, etc.). Furthermore, by putting the passengers in the car control, it allows them to make the decision that leads to the consequences they feel the 'best' about, not simply forcing them into an outcome they have no control over."*

Here the student refers to the legal issues associated with the liability to the company instead of any moral underpinnings.

6.2 Analysis 2: When students chose to swerve 'RIGHT' (towards the motorcyclist with helmet) on Condition 1: SELF *(31/82)*

There were thirty-one students who chose to go right in condition 1. Their reasoning was primarily based on the fact that the motorcyclist on the right did wear a helmet which, according to some increased his chances of being saved in a case of collision. The reasoning here was to minimize the harm while saving the passengers in the car. This is an example of *consequentialist* decision making which "holds that the consequences of one's conduct are the ultimate basis for any judgment about the rightness or wrongness of that conduct" [17]. An example of a student's response is the following:

> *"I would steer right and hit the motorcyclist with the helmet on. ....Although I would be hitting a motorcyclist with a Helmet on it is the only option with the possibility of not completely killing the motorcyclist. Of course I wouldn't just let my car ram into the truck, I would by all means try to survive the accident weather or not it's at the expense of another live-simply survival"*

This response again confirms the consequentialist nature of student's decision. However, when the same students were asked to decide in the algorithmic context, fifteen out of thirty-one explicitly indicated an adjustment to their decision. Eleven out of these fifteen students indicated that algorithms should not make any decision. This suggests that while one-third of the total students chose to go right in condition 1, approximately half of them (fifteen out of thirty-one) did not think that an algorithm should make any decision which is an adjustment from consequentialist reasoning to non-consequentialist reasoning. The following is an example of a student's response in an algorithmic context given that they chose to swerve right in condition 1:

> *"I will propose the algorithm to tell the car [that it] doesn't take any action itself and leave the decision to the passengers in the car. The passengers can take the control at that time and make decision to drive either way under that situation. I said before AI can't decide incidents with ethics issues like this one. And also if we make the car to decide, any damage caused by this accident will be charged by our company and the breaker factory maybe. Machines can't be designed to decide which decision will be the best according to different witness and victims with different outcomes and different probability to avoid the accident"*

So, we see that students both in earlier analysis and in this analysis increasingly share the opinion that there are legal issues associated with algorithms making such decisions. Thus they indicate a preference for humans deciding the course of action rather than standardizing it through an algorithm.

6.3 Analysis 3: When students chose to go 'STRAIGHT' (towards the truck) on Condition 1: SELF *(37/82)*

Thirty-seven students out of eighty-two chose to go straight (by not taking any action) in condition 1. The majority students decided to go straight primarily because they relied on the safety features of the car (like airbags or physical properties of the car to survive a collision) which they thought would be sufficient to protect them. Furthermore, by going straight, they would also not hurt anyone else. This is again an example of *consequentialist* decision-making by the students. An example of a student's response is:

> *"I would do nothing and keep going straight. The individuals in the motor cycle would not be able to survive the hit from a car at 70 mph whether or not they have a helmet. However, being inside the car, there's a higher chance that I would survive crashing into the truck as a result of the protective equipment in the car like airbags and seat belt"*

This response confirms the consequentialist nature of the student's decision. However, some students also expressed that since they were responsible for the condition, they should not put someone else's life in danger. Some students also expressed that because they think all lives are equal, they would not prioritize their own lives over others. They indicated that they should not play with the chance or the luck of any individual which can be classified as the luck-egalitarianism viewpoint. *Luck-egalitarianism* viewpoint "expresses the intuition that it is a bad thing for some people to be worse off than others through no fault of their own" [18]. The following are some of the students' responses who expressed this viewpoint:

> *"I would hit the truck myself. There is no need to take a chance at hurting or killing someone else when it is something I have caused myself. Taking someone elses life to save my own is the definition of selfish."*

> *"I don't believe it is within my own right to say my life is more valuable than either motorcyclist (regardless of age). There is a high possibility that I can kill either motorcyclist if I steer into them (with the motorcyclist with the helmet having a slightly higher chance of survival). Granted, if I did choose to hit either motorcyclist, there is a higher chance of survival for myself since I am in a larger vehicle. It is better to do nothing and let my car hit the truck with the possibility that I may die or become severely injured since the truck is larger. This way, I do not have to make a decision of possibly taking a person's life just for the sake of my own survival. My quality of life living with the decision of killing an innocent human/bystander would be worse than not living at all."*

So, while the students' decision to go straight (or do nothing) is same, it stems from different ethical reasoning frameworks.

When the same students were asked to decide in the algorithmic context, eighteen out of thirty-seven explicitly indicated an adjustment to their decision, and fourteen out of these eighteen students indicated that algorithms should not make this decision. This suggests that approximately more than one-third of the students who chose to go straight in condition 1 did not think that algorithm should make any decision in condition 2. The following example are some student's responses:

> *"I would probably program the car to not take any action. I don't think a computer should have the capability to decide which life is more valuable, or the capability*

> *to knowingly harm a human. Therefore when scenarios like this arise, I think the best course of action would to have the computer shut down and revert back to manual driving to allow the human driver to make the decision."*
>
> *"I think that if possible I would write a program checking the brakes before they went on the road or at least check the status of the breaks to see if it is safe to drive. But in this scenario and I had to choose I could never code a car to kill someone else. I would either do nothing or let the person in the drivers seat take over and make that decision for themselves."*

The above findings present major changes in students' decision-making in the two conditions and the potential philosophical frameworks from which students draw their reasoning from to explain their decisions.

7 DISCUSSION

In the analysis presented above, we can see that firstly in condition 1 a majority of students either choose to go straight primarily because of the reliability of the car's safety features while not endangering any motorcyclist's life. Or, students decide to go to the right to minimize the loss of lives as they save themselves while potentially saving the motorcyclist as he is wearing a helmet. However, when the same groups of students were presented with the condition 2, regardless of their decision in condition 1, one-third to half of the students indicated that algorithms should not be tasked with life-and-death decisions, even if the algorithm's decision would comport with what the student believed a human driver should do. This is interesting because while the algorithmic case does change the situation by bringing factors like the legal liability of a company or responsibility of the car owner, there is a large number of students who think that, fundamentally, algorithms should not make such decisions.

Figure 6 and 7 shows the overall reasoning based on individual decisions of students in the two conditions. We can see that while the presence of helmet and safety features of car dominantly influence decision-making in condition-1, algorithms should not make the decisions become dominant in condition-2.

Now, in the given scenario, it was a question of life and death. However, will these students reason in the same way when it comes to algorithms making decisions on predictive policing or university admissions? While cases like these may not be directly life-death questions, they certainly determine life-changing outcomes for an individual. In future research, it would be interesting to study students' reasoning on the use of algorithms in such cases.

8 LIMITATIONS

The current study attempts to bring light to the mindset of students undertaking CS education to the situations which require deep ethical understanding as AI gains more ground in our day to day life. This study shows the common themes as well as dissimilarities in reasoning behind their decisions. Limitations of the current study are the small sample and sampling of students from one institution. Though from the statistical viewpoint and external

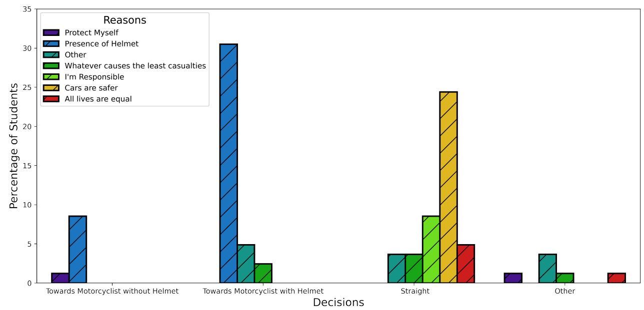

Figure 6: Students' reasons grouped by decisions on Condition 1: SELF when they are themselves driving the car

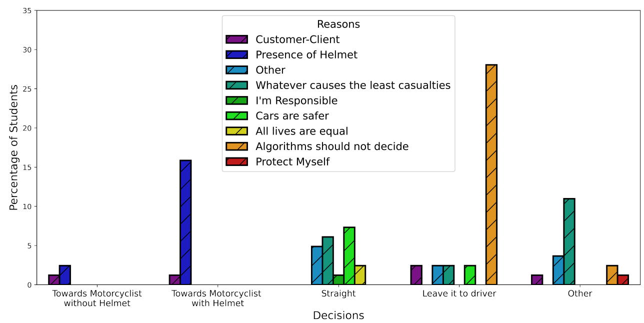

Figure 7: Students' reasons grouped by decisions on Condition 2: ALGORITHM when they are writing an algorithm

validity these two things are important, the purpose of the current work was to exploratory investigate the subjective nature of ethical norms students hold in their mind while they design an autonomous vehicles and set an agenda for future research programs for ethics education among CS-Ed community.

9 CONCLUSION

As the CS-Ed community explores new ways to cultivate ethical-CS understanding in undergraduate students, it is essential to recognize the subject's inherent complexity, especially added by the rapid advancements at the intersection of technology, economics, social and political circumstances. While much work has been done in engineering ethics, there is a potential to explore the realms of ethical algorithms for informing the undergraduate curriculum. In this paper, we present the study results that analyzed how undergraduate engineering students enrolled in a CS1 course reason about ethical and algorithmic decision-making processes.

The goal of conducting such an analysis was to understand students' default reasoning patterns and their philosophical underpinnings, which can help us understand how students generally think about ethical decision-making processes. And secondly, to create a curriculum that considers these perspectives for it to be engaging and meaningful. We believe that this will facilitate a more human-centric participatory approach where students' viewpoints are actively considered in informing and guiding the CS-Ethics curriculum.

ACKNOWLEDGMENTS

The authors would like to thank Caleb Wong for his help in preliminary analysis and Duncan Purves for his feedback and suggestions. We appreciate the detailed reviews from the SIGCSE reviewers.

REFERENCES

[1] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. The moral machine experiment. *Nature* 563, 7729 (2018), 59–64.

[2] Christopher W Bauman, A Peter McGraw, Daniel M Bartels, and Caleb Warren. 2014. Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral psychology. *Social and Personality Psychology Compass* 8, 9 (2014), 536–554.

[3] Elettra Bietti. 2020. From Ethics Washing to Ethics Bashing: A View on Tech Ethics from within Moral Philosophy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 210–219. https://doi.org/10.1145/3351095.3372860

[4] Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2016. The social dilemma of autonomous vehicles. *Science* 352, 6293 (2016), 1573–1576. https://doi.org/10.1126/science.aaf2654

[5] 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* (St. Louis, Missouri, USA) *(SIGCSE '05)*. Association for Computing Machinery, New York, NY, USA, 347–351. https://doi.org/10.1145/1047344.1047464

[6] Stacy A. Doore, Casey Fiesler, Michael S. Kirkpatrick, Evan Peck, and Mehran Sahami. 2020. Assignments That Blend Ethics and Technology. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (Portland, OR, USA) *(SIGCSE '20)*. Association for Computing Machinery, New York, NY, USA, 475–476. https://doi.org/10.1145/3328778.3366994

[7] Seamus F Freyne, James P Abulencia, and Powell Draper. 2010. First Year Engineering Students' Perceptions Of Contemporary Ethical Issues. In *2010 Annual Conference & Exposition*. ASEE Conferences, Louisville, Kentucky. https://peer.asee.org/16419

[8] Seamus F Freyne, James P Abulencia, and Powell Draper. 2011. Engineering and Nonengineering Students' Perceptions of Contemporary Ethical Issues. In *2011 ASEE Annual Conference & Exposition*. ASEE Conferences, Vancouver, BC. https://peer.asee.org/17857

[9] Heidi Furey and Fred Martin. 2019. AI Education Matters: A Modular Approach to AI Ethics Education. *AI Matters* 4, 4 (Jan. 2019), 13–15. https://doi.org/10.1145/3299758.3299764

[10] Barbara J. Grosz, David Gray Grant, Kate Vredenburgh, 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

[11] Brittney Hope Jimerson, Eui Hyun Park, Vinod K Lohani, and Steven M. Culver. 2013. Enhancing Engineering Ethics Curriculum by Analyzing Students' Perception. In *2013 ASEE Annual Conference & Exposition*. ASEE Conferences, Atlanta, Georgia. https://peer.asee.org/19544

[12] Christoph Lutz and Aurelia Tamò. 2015. RoboCode-Ethicists: Privacy-Friendly Robots, an Ethical Responsibility of Engineers?. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Ethics in Networked Systems Research* (London, United Kingdom) *(NS Ethics '15)*. Association for Computing Machinery, New York, NY, USA, 27–28. https://doi.org/10.1145/2793013.2793022

[13] C. Dianne Martin. 2011. Reasoning with Ethics. *ACM Inroads* 2, 1 (Feb. 2011), 8–9. https://doi.org/10.1145/1929887.1929889

[14] Elizabeth Milonas. 2020. How Do We 'Raise' Ethically Minded Computer Students?. In *2020 ASEE Virtual Annual Conference Content Access*. ASEE Conferences, Virtual On line. https://peer.asee.org/34730

[15] Ali Salehnia and Shieva Salehnia. 2016. Ethical Issues in Computing: Student Perceptions Survey. In *2016 ASEE Annual Conference & Exposition*. ASEE Conferences, New Orleans, Louisiana. https://peer.asee.org/26740

[16] 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* (Baltimore, Maryland, USA) *(SIGCSE '18)*. Association for Computing Machinery, New York, NY, USA, 940–945. https://doi.org/10.1145/3159450.3159573

[17] Wikipedia contributors. 2021. Consequentialism — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Consequentialism&oldid=1037766712. [Online; accessed 14-August-2021].

[18] Wikipedia contributors. 2021. Luck egalitarianism — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Luck_egalitarianism&oldid=1012462990. [Online; accessed 14-August-2021].

[19] Wikipedia contributors. 2021. Retributive justice — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Retributive_justice&oldid=1038190163. [Online; accessed 14-August-2021].