

RETHINKING DATA SCIENCE PEDAGOGY WITH EMBEDDED ETHICAL CONSIDERATIONS

Vandana P. Janeja, Maria Sanchez

University of Maryland, Baltimore County

Abstract

The focus of this paper is to present a tool to meet the need of developing ethical critical thinking in data science curricula for undergraduate students. New data science methods impact societies, communities directly or indirectly when dealing with open and other real-world datasets. In particular, for data science there is a need to develop ethical critical thinking while analyzing the data. In the knowledge discovery process there are many opportunities for ethical decision making that a data scientist can evaluate throughout the entire lifecycle of the data to do no harm. To address these concerns within a learning environment focused on skilled workforce development we first introduce a novel ethical data lifecycle framework and then propose a vehicle for implementation through a short term module that can be embedded into a fast paced data science course. The objective of the module is to increase the ethical thinking of students when analyzing data. Pre and post surveys were conducted across each of the two semesters to evaluate students' attitudes towards ethical thinking. The analysis of the survey results suggests that the objective was achieved based on a positive shift toward agreement with statements related to the importance of ethical thinking.

Keywords: ethical critical thinking, data science, data life cycle

1 INTRODUCTION

Conversations regarding data science and ethical concerns have become more prevalent over the last several years [1, 2, 15]. Data Science curricula often focuses on the fast paced delivery of skills that are much needed in the workforce. At the same time there is a need for slowing down the delivery to allow for important ethical concerns to be addressed. There are several serious questions surrounding data analysis such as, who owns the data that is collected? Who regulates the use of personal information and determines responsibility for the mishandling of data? And do users of technology understand what personal data is being collected and what are the possible repercussions related to the misuse of that data?

Higher education institutions are trying to address these concerns by focusing on educating technology creators on different fronts and at the same time making an impact on the community [15]. The analysis of big data is a critical component in the development of all aspects of modern technology, and it is critical that data science students have a proper understanding of the associated ethical issues in this domain.

A large and much appreciated resource for the data science community is the collection of ethics curricula [6] and the paper presented in SIGCSE analyzing these curricula [5]. To undertake the development of an embedded ethics module described in this paper, we performed our own analysis of several of the syllabi extracted from the list [6]. We analyzed over 100 syllabi from the open spreadsheet and identified trends present in the classroom (such as [9-13]. The most common topics encountered in the analyzed courses include privacy, social media, policies governing technology such as General Data Protection Regulation (GDPR), intellectual property, surveillance, and self-driving cars [13]. Many courses focused on access to individual personal data through Big Data Giants such as Facebook and Google. Courses on ethics in computer science departments primarily focused on fairness in machine learning and natural language processing and bias introduced through algorithms. In terms of assignments, we observed that most of the courses followed a pattern in which there were reading assignments to be completed before the class from books, web articles from reputed sources such as New York Times, Washington Post, etc. [7]. Despite the variety of topics encountered in the courses, there is very limited work involving the use of datasets or a collection of datasets to help students understand the risks involved in using real-world data and gain hands-on experience.

While analyzing the courses, we observed that there were two different approaches to instruction of ethics in data science. The first approach included the use of ethics modules for a few weeks within the broader course in data science. The other approach consisted of standalone ethics courses where data science was one of the topics discussed. The ethics modules normally focused on topics such as AI, Machine Learning, IoT and Big Data and use case studies related to intellectual property, privacy security, surveillance of internet information and the effects of the mishandling of data. Standalone ethics courses have a more comprehensive coverage of ethics affecting society in different ways. When discussing ethical consideration in data science, the most common topics included impact of social media, AI, Machine Learning and Bias. These courses typically have more extensive reading assignments such as books [7, 8, 10] and articles and more in depth case studies. Learning from this landscape study we went with a module approach to be able to embed it into existing data science classes so that the ethics discussions are in context of the data lifecycle.

Our contribution in this paper are twofold: (1) We propose a novel framework of an ethical data life cycle which captures the nuances of ethical considerations throughout the data life cycle and (2) We utilize this framework in a module that can be embedded into a data science class along with other discussions around the theoretical underpinnings of ethics and the practical applications of the data life cycle, through discussions of case studies.

Our approach offers a versatile educational tool for delivery in the form of a virtual module. The framework was first developed in 2019 [3] and the module has since been embedded in multiple data intensive courses (such as data science and cyber informatics) offered to students from multiple technical disciplines. Results evaluating the impacts of the module on students' perceptions towards ethical considerations in data science, based on a survey administered in the course for two semesters, are presented and discussed in this paper.

The rest of the paper is organized as follows: section 2 discusses our approach including the framework and module. Section 3 discusses the results of evaluating the module outcomes across two semesters of implementation, where the module was embedded into a foundational data science class. We finally conclude and share future directions in section 4.

2 APPROACH

In our approach we present the core theoretical framework around which we have created a pedagogical intervention to support developing ethical critical thinking in students. This is done using an ethical data life cycle. We adopt this framework into a module that we have found to be generalizable in data science classes or classes which cover data analytics topics.

2.1 Framework for Ethical thinking in Data Science

Ethics generally works on the principles of not only do-no-harm but also to take action to lead to fair results. Although research protocols to protect human beings have been in place for a while now, new data science methods relying on digital traces that people leave, make it less clear where the human beings are in data science work. Harm is not only individual, nor does it necessarily rely on intervening in someone's life directly, particularly given the availability of information on the internet and the practice of connecting datasets. As we deal with more and more open datasets the potential to do harm to individuals and even communities increases.

In particular, for data science there is a need to develop ethical critical thinking while analyzing the data. In data driven decision making AI systems, there are several checkpoints where a system follows a pathway based on the choices that the creator of that system made or the user of the system made. It is important to know what those choices are and how they can impact people's lives and livelihoods. It is also important to explain the outcomes from such systems and ensure they are not disproportionately impacting certain individuals. Just like we would carefully hire an employee, supervise them, give them advice in making critical decisions, the sensible extension would be that we should do the same for data driven decision making systems which make decisions based on bulk datasets.

To alleviate this situation, throughout the entire lifecycle of the data in the knowledge discovery process[14] there are many opportunities for ethical decision making that a data scientist can evaluate to do no harm.

Here the harms are not only through identification of Personally Identifiable Information (PII), but also other types of data that may not necessarily fall under the direct purview of traditional IRB protections. These other types of data may indirectly lead to identifying human behaviors that should fall under privacy concerns. In addition, ethics may not always equate to privacy and there could be other types of data, which may not directly be identifiable but still have to be carefully handled. There may be yet other types of data, which if

handled without an ethical viewpoint may have adverse impacts for the society as a whole. Data science is generally presented in the data life cycle pioneered through the knowledge discovery data mining life cycle[14]. Using this data life cycle as a backdrop, ethics cannot be tacked onto one part of the data life cycle but across the data life cycle infused into the process of discovering the patterns in the data. An ethical data lifecycle, first described in 2019 [3], is proposed in this project and depicted in Figure 1.

Each part of this lifecycle must be observed and assessed with ethical considerations which may emerge in the form of questions to ask as these stages are framed in a project. This process has to start at data collection, beginning with the considerations of diversity of populations and privacy of individuals protected. During data integration ethical considerations play a key role in deciding whether data should be reused or combined with other datasets. During data preprocessing questions should be carefully evaluated, such as: Should missing values be filled with averages from a broader population? Should anomalies be excluded from the data or included for a more careful assessment of the rare nuggets in minority data distributions? Bias should be carefully avoided during data selection for the task being evaluated. Data pre-preprocessing should consider sampling strategies, selection bias which can again creep in at this stage due to groups selected while preprocessing the data such as through geospatial and temporal context in selecting the data and forming the right groups. In the next stage, in which data is processed through pattern discovery algorithms and thresholds, threshold choice can impact what patterns are discovered or excluded. In addition, aspects of the provenance of the algorithm and the data, the reproducibility of the results, and many other decision points during this phase have far-reaching ethical impacts. Once the patterns are discovered, it is essential that pattern evaluation include checks for detecting implicit bias and tested against ground truths to establish not only accuracy but also veracity of the data and the results. Here veracity refers to trustworthiness of the data and results. If the data is imprecise and not representative, the results also will reflect this imprecision and non-representativeness. In addition to the ethical concerns at each step of the data life cycle, a choice at one stage in the data life cycle can impact the following stages and result in compounded ethical issues. For example, if data collection is not representative of the population, then even if the algorithm is robust to bias, the patterns detected will be biased. So, in essence, the model is only as good as the data and the patterns are only as good as the model. Knowledge gained from this cycle can go back into the cycle in a feedback loop, informing the data collection phase again. Looking at the entire data life cycle from an ethical lens is key to any successful and societally impactful project. Such a thinking is important to discuss in a data science class. The module we propose in this paper embeds such an ethical data life cycle along with theoretical underpinnings of ethics.




Figure 1: Ethical Data Lifecycle

2.2 Module Structure

In the data analytics classes including data science and cyber informatics, students are learning the skills they needed for the workforce in data science. Students make design choices in their algorithms with thresholds and branching of algorithms without having a deeper thought around the ethical framing of it. So, this module was designed to embed ethical discussions in a fast paced skill driven classroom in a short amount of time. The Ethical thinking in Data Science (Ethics-DS) module consists of two lectures. The first lecture focuses on theoretical concepts of ethics such as ethics vs. morality, ethical theories, code of con-

duct, etc. and it is delivered by a domain expert. In the second lecture the Ethical Data Lifecycle highlighted in Section 2.1 is introduced. Through real-world case studies, the lifecycle is used as a mechanism to encourage students to ask ethical questions and evaluate the appropriate ethical considerations associated with the scenarios presented. Pre and post surveys were conducted to assess the students' perceptions about ethical considerations in data science. Students were also asked to complete reflection questions associated with the case studies discussed.

2.3 Ethical Data Life Cycle

There is one concept that every data science student learns which is this data lifecycle. We created this part of the module adapting the data life cycle as described earlier. The idea here is that using this concept of the life cycle the students learn to ask questions for every step of that lifecycle and this makes students think of algorithmic design choices that can create an impact on associated individuals whose data is being processed. We then talk about case studies that use this lens of the data lifecycle and make decision choices. Example case studies included current topics (for example COVID related contact tracing). Just with these two lectures we saw that students started thinking more deliberately, actionably around the choices that they were making in their projects and assignments.

2.4 Adoption

The module includes voice over lectures that can be embedded in the curricula. We also provide a courseware module with quizzes that can be linked to grading in the course shell. We have recently also added Jupyter notebook examples (such as K-Nearest Neighbors classification) where we take an existing data science notebook and show variations in algorithmic choices and how that impacts the accuracy of the classifier. The course module is available through GitHub [4]. This module can be embedded in any data science class and also in classes that discuss data analytics topics. We have experimented with this module in foundations of data science class and also a cyber informatics class and have found this to be fairly generalizable where a data life cycle is discussed.

3 RESULTS

3.1 Data and population

We present results from two offerings of the foundations of data science class where we embedded the ethics module. In Spring 2020 we had 12 students in the class and in Spring 2021 we had 18 students from a variety of majors. The rate of response from Spring 2020 was 91.7% while the rate of response in Spring 2021 was 77.8% Table 1 shows the distribution of the respondents to the survey in both semesters.

3.2 Comparison over two semesters

A series of statements were presented to the students to measure the effect of the course on their attitudes related to the ethical considerations in the analysis of big data. Particularly we included questions around (1) perceptions towards ethical considerations (Figures 2, 3, 4) and (2) actions (Figures 5, 6, 7) they may need to take as a result of their understanding of ethical issues. The following are a sampling of the results.

Figures 2 (a) and (b) present survey results for the statement "Any data science project must have discussion of ethical implications" for the 2020 and 2021 courses, respectively.

Table 1. Distribution of respondents of survey

Semester	# of responses	Majors
Spring 2020	11 (6 Females; 5 Males)	3 majors (IS, Comp Sci, BTA)
Spring 2021	14 (3 Females; 11 Males)	6 majors (IS, Pre-Comp Sci, BTA, Math, Psych, STAT)

In the 2020 offering there was a slight decrease in the Somewhat disagree/Neither agree nor disagree/Somewhat agree categories and an increase of the number of students that agree with the statement

over the duration of the course. This trend was more pronounced in the 2021 offering as at the end of the class all students have shifted to the Agree/Strongly agree categories. This reflects on their understanding of the importance of ethical considerations. Figure 3(a) and (b) illustrate the survey results regarding the statement “Understanding Data Science and ethical implications will help solve important problems in a more just and fair manner”. A similar trend to the previous figure is observed where the students’ attitudes tend to be more in agreement with the statement after going through the module. Figure 4 (a) and (b) show the survey results related to the statement “A person whose data is being collected should be the one to decide what happens with their data”.




Figure 2. Perception - importance of ethical discussions (a) 2020, (b) 2021




Figure 3. Perception - ethical understanding for just and fair problem solving (a) 2020, (b) 2021



Figure 4. Perception 2020 Importance of Data ownership (a) 2020, (b) 2021

We next discuss student responses towards action statements. These statements are designed to elicit responses to what the students would need to do in their day to day lives as data scientists. We anticipate that it takes time for students to become well versed with the action they need to take. In these figures, although there is still a positive trend towards agreement (especially in figures 5 (b) and 7(b)) with the action statements, there are students that still tend towards disagreement (figure 6(a)) or neutrality (figure 5(a), 6(b)) in response to statements at the end of the module. This may indicate that the module is more effective in raising awareness of ethical considerations, but it may not be enough to translate that awareness to confidence in taking action to resolve ethical issues. This is not surprising since we anticipate that especially in the undergraduate population it takes time to go from awareness to action. This also makes sense to support the argument that one model may not fit all ethics education in data science.

Figure 5 shows the survey results related to the statement “A person whose data is being collected should be the one to decide what happens with their data”. Figures 5 (a) and (b) show that the module seems to influence the students’ attitude on the question of ownership of data with the higher percentages of students responding in agreement with the statement at the end of the module than at the beginning.




Figure 5. Action 2020-Types of questions to ask(a) 2020, (b) 2021




Figure 6. Action -dealing with impact in real world projects (a) 2020, (b) 2021




Figure 7. Action -Looking up information on code of conduct (a) 2020, (b) 2021

Overall, the ethics module has had a positive effect on the attitude of the students towards the importance of ethical considerations in the analysis of data. It is also important to note that, in general, students in the 2021 offering tended to start the module with a higher tendency towards neutrality and agreement to the action statements (Figures 5(b), 6(b), 7(b)) than those students in the 2020 offering. Essentially the responses are somewhat right skewed. A possible explanation for this occurrence could be the effect of the increased exposure to data associated with the pandemic which may have brought in students at a higher level of awareness. However, this is not necessarily the case for the perception statements. It is our hypothesis that perceptions are complex to understand but easier to establish, whereas actions while they take time to establish are concrete and easier to actualize. The goal of the module was to increase the ethical thinking of

students when analyzing data. The survey results demonstrate that the goal was achieved as there was a general increase of the percentage of students that agree with the statements presented about ethical thinking.

4 CONCLUSIONS AND FUTURE WORK

In order to educate future data science students in ethical issues, we proposed an approach that consists of a theoretical framework and a virtual module for its implementation. The theoretical framework named Ethical Data Lifecycle encourages students to ask the appropriate ethical questions throughout each step of the data life cycle process. The Ethics-DS virtual module is used as a vehicle for the implementation of the framework as well as the instruction of ethical theories applied to data science real life scenarios and can be embedded in a data science class. Results from pre and post surveys during two semesters demonstrate that the module is effective in increasing the ethical thinking of data science students. During the delivery of the module in the two semesters, students were asked to complete reflection exercises on the ethical scenarios discussed in class. In our future work, we wish to analyze the students' responses to identify patterns that could provide insight into the depth of their understanding of ethical issues in data science and in particular action oriented tasks.

5 REFERENCES

- [1] Wing, J. M., Janeja, V. P., Kloefkorn, T., Erickson, L. C. (2018). Data Science Leadership Summit: Summary Report.
- [2] Executive Summary, Including Ethics in Data Science Pedagogy (NSF-EDSP- 2019), June 17-18, 2019, Alexandria, VA, UMBC: Shimei Pan, Jimmy Foulds, Susan Sterett, Vandana Janeja, UCB: Cathryn Carson , Georgetown: Lisa Singh, UW: Bill Howe, <https://sites.google.com/umbc.edu/edsp19/resources?authuser=0>
- [3] V. P. Janeja, Do No Harm: An Ethical Data Life Cycle, AAAS Sci on the Fly, April 2019.
- [4] Ethics Module GitHub, <https://github.com/MultiDataLab/EthicalDataScience>
- [5] Fiesler, C., Garrett, N., Beard, N. (2020, February). What do We teach when We teach tech ethics? A syllabi analysis. In Proceedings of the 51st ACM Technical Symposium on Computer Science Education (pp. 289-295).
- [6] Tech Ethics Curriculum , Google Sheets ,<https://bit.ly/3CCAOXT>, Last Accessed: 28th July, 2021
- [7] Instructor: Yulia Tsvetkov, Course Title: Computational Ethics for NLP, Department: Language Technologies Institute, University: Carnegie Mellon University, URL: <http://demo.clab.cs.cmu.edu/ethicalnlp/> Last Accessed : 28th January, 2021
- [8] Instructor: Prof. L. Felipe Perrone, Course Title: Computers and Society, Department: Computer Science, University: Bucknell University, URL: <http://www.eg.bucknell.edu/cs240/syllabus/> , Last Accessed: 28th January, 2021.
- [9] Instructor: Beth Harnick-Shapiro M.A. , Course Title: Professional Ethics in Computing , Department: Computer Science, University: California State University Fullerton, URL: <https://goo.gl/pu1hda> , Last Accessed: 28th January, 2021.
- [10] Instructor: Katie Shilton , Course Title: Information Ethics , Department: College of information Studies, University: University of Maryland College Park, URL: kshilton@umd.edu, Last Accessed: 28th January, 2021.
- [11] Instructors: Weitzner, Fischer, Edelman , Course Title: Foundations of Internet Policy , University: Massachusetts Institute of Technology, Article Title: Exercise - Class 12 - Nov 30, 2017 - Digital Platforms, URL: <https://bit.ly/3fURLwl>, Last Accessed: 28th January, 2021
- [12] Barocas, S., Boyd, D. (2017). Engaging the ethics of data science in practice. Communications of the ACM, 60(11), 23-25.
- [13] Instructor: Phillip Rogaway, Course Title: Ethics in an Age of Technology, Department: Computer Science, University: UC Davis, URL: <http://www.cs.ucdavis.edu/blog/ecs-188-ethics-age-technology/>, Last Accessed: 28th January, 2021
- [14] Mariscal, G., Marban, O., Fernandez, C. (2010). A survey of data mining and knowledge discovery process models and methodologies. The Knowledge Engineering Review, 25(2), 137-166.
- [15] Engineering National Academies of Sciences and Medicine. 2018. Data Science for Undergraduates: Opportunities and Options. The National Academies Press, Washington, DC