



Data Science in a Mathematics Classroom: Lessons on AI Fairness

Berenice Sotelo

Northwest Early College High School
El Paso, USA

bsalazar@canutillo-isd.org

Kirsten Wieseman

Northwest Early College High School
El Paso, USA

kwieseman@canutillo-isd.org

Saeid Tizpaz-Niari

University of Texas at El Paso
El Paso, USA

saeid@utep.edu

ABSTRACT

Mathematics forms the foundation of computer science. By applying mathematical concepts to computer science, students can gain a deeper appreciation for the significance of mathematics and become more engaged in the learning process. This paper explores how integrating data science and AI fairness concepts into two distinct mathematics courses—Discrete Math and Algebra 1—enhances student learning and engagement.

In Discrete Math, an upper-level course focusing on discrete mathematical structures, data science applications provide concrete examples of abstract concepts. For instance, recommendation systems and their fairness can be used to motivate graph theory. Similarly, in Algebra 1, a foundational course for future mathematics study, AI fairness problems offer real-world contexts for algebraic concepts. Students can explore linear equations and inequalities by examining AI decision-making processes and their potential biases.

By incorporating these data science and AI fairness elements, we transform potentially abstract or mundane mathematical activities into meaningful, relevant learning experiences. This approach not only facilitates a deeper understanding of mathematical concepts but also introduces students to critical ethical considerations in technology.

1 INTRODUCTION

AI has been increasingly integrated into important software infrastructure to support decision-making processes with significant socio-economic and legal implications. By leveraging data, AI-driven software can perform complex tasks over high-dimensional input data. Such software has also made substantial intruder into high-stake domains such as software deciding on recidivism, software determining whether DNA was in the evidence sample, and software deciding whether to audit a given taxpayer.

These rapid adoptions of AI-driven solutions, however, pose major threats to the safety of automated decision-support software. Unfortunately, there are plenty of fairness bugs in AI-driven software, especially as the epistemic divide between these systems and field experts becomes apparent in edge cases. For instance, COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), a risk assessment software in criminal justice, has shown a troubling bias against Black defendants by falsely predicting a higher risk of re-offending for them than others [3]. Since AI learns its logic from prior experience and dataset, it is not surprising that they encode historical and present biases due to displacement, exclusion, segregation, and injustice. The resulting software may particularly disadvantage minorities and protected groups and be found non-compliant with law such as the US Civil Rights Act.




Figure 1: FairLay-ML Toolkit.

In this paper, we explore how to integrate data science and AI fairness concepts into two distinct mathematics courses—Discrete Math and Algebra 1—enhances student learning and engagement. We ask the following research question: how can we use case studies like COMPAS [6] to enhance students’ learning and engagements in two distinct mathematics courses—Discrete Math and Algebra 1?

Discrete Math, an upper-level course focusing on mathematical structures that are fundamentally discrete, offers numerous opportunities to explore AI and data science concepts. Topics such as graph theory, combinatorics, and logic directly apply to various aspects of AI systems. For instance, graph theory can be used to understand social network analysis in AI applications, while combinatorics is crucial in data mining algorithms. Algebra 1, a foundational course for future mathematics study, can leverage AI fairness problems to provide real-world contexts for algebraic concepts. Students can explore linear equations and inequalities by examining AI decision-making processes and their potential biases, making abstract concepts more concrete and relevant.

In the future, we want to implement the skills we acquired in our classrooms. In doing so, we plan to leverage Parfait-ML [1, 2, 4, 5, 7] GUI app, FairLay-ML [8, 9], available at <https://fairlayml-v2.streamlit.app/>. Figure 1 shows an overview of the toolkit.

2 DISCRETE MATHEMATICS AND AI FAIRNESS

Discrete mathematics is the study of mathematical structures that can be considered discrete (not continuous). Concepts and notations from this field are useful in studying and describing objects and problems in computer science.

Logic Activities. The first activity will introduce logic by using conditional statements. As a class, we will write conditional statements about simple tasks. For example, we will write about what to do if you are single, have diploma, etc. Then, students will work individually to write their own conditional statements. In the second activity, students will create an If/Then story with at least three different decisions using Google Slides. The presentation must be creative, and the decisions need to be linked to a different slide. Students will then present to the class. In the third activity, students will be given nine different problems to choose from. They will be using Python to write code to fulfill the requirements. The program they will be using already has the function defined. Students need to write if, else if, and else statements. The program checks your code using different scenarios. If the code works the scenario will turn green, and if it doesn't work it turns red. In the case that it turns red, students need to modify their code until every scenario is green. In the last activity of logic, students will be doing a hands-on activity to program a Sphero so that it goes through a maze. Sphero is a programmable robot that can be programmed using blocks or JavaScript, depending on the students' abilities.

Data Analysis. During the first day of this unit, will present AI Fairness topic. Then students will have a discussion over artificial intelligence, machine learning, and machine bias. In the next activity, students will use what they learned in the logic activities. For example, students use conditional statements to infer decisions based on relevant factors. However, this could be problematic like when we compared sex, age, and race. We will walk students through the FairLay-ML toolkit where they can investigate fairness of conditional programs and visualize different decisions. In activity three, students will design a Google form to collect data from Northwest Early College High School student population. They will need to think about what data they want to collect and how they are they going to analyze it. In activity four, students will use the Google Colab and/or FairLay-ML to analyze the. Then, they will create a presentation to share their findings.

Discrete TEKS: 1.A, 1.B, 1.C, 1.D, 1.E

3 ALGEBRA 1 AND AI FAIRNESS

Algebra 1 is the study of linear, quadratic, and exponential functions and their related transformations, equations, and associated solutions. Students will work with functions and their solutions in real-world situations while using technology to collect, explore data, and analyze statistical relationships.

Slope Activities. In our research, through the analysis of slope, we were able to evaluate the fairness of COMPAS. Similarly, students will be able to evaluate types of trends and correlations using slope. I will begin with introducing coding through two activities which incorporate slope. The first activity is Slope Art with Python, where students will create a piece of artwork by coding lines and polygons. The second activity is the Slope Formula with Python. This activity consists of coding the slope formula and finding various types of slopes, positive, negative, zero, and undefined, if given two points.

Scatterplot Activity. Machine learning uses patterns in data to draw conclusions and create predictions when approaching new data. Scatterplots, made up of a scattered group of points, are much

the same as the data points in a data set used for machine learning. We are able to analyze scatterplots for various types of trends or correlations, similarly to our analysis of the COMPAS dataset. I will allow my students to choose their favored data collection instrument such as: survey, experiment, physical measurements, or reputable website or article. The students will create a depiction of their dataset as a scatterplot on a poster board. Students will check for the type of correlation present, positive, negative, or no correlation, in their scatterplot.

Linear Regression Activity. Since a variant of linear regression models, known as logistic regression, used to classify defendants into low, medium, and high risks in the COMPAS software, we plan to also motivate students to infer linear regression models and understand their socio-fairness implications. We will encourage students to use applications with deep socio-economic implications such as COMPAS, Adult Census Income Dataset, Credit Dataset, etc. By measuring fairness of linear regression models with metrics like true positive rate differences between white and Hispanic, students can quantify the amount of discrimination.

Algebra 1 TEKS: 3.A, 4.A, 4.B, 4.C

ACKNOWLEDGMENTS. We would like to thank our mentor Leo Mercado for all his patience and help. We would like to thank Drs. Ceberio, Robertson, and Tosh for the inspiration and optimism. This project is supported by NSF under the grant CNS-2206982. Tizpaz-Niari is partially supported by NSF under grants CNS-2230060 and CCF-2317206.

REFERENCES

- [1] Vishnu Asutosh Dasu, Ashish Kumar, Saeid Tizpaz-Niari, and Gang Tan. 2024. NeuFair: Neural Network Fairness Repair with Dropout. *The ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA'24)* (2024). <https://doi.org/10.48550/ARXIV.2407.04268> arXiv:2407.04268
- [2] Salvador Robles Herrera, Varya Monjezi, Vladik Kreinovich, Ashutosh Trivedi, and Saeid Tizpaz-Niari. 2024. Predicting Fairness of ML Software Configurations. In *Proceedings of the 20th International Conference on Predictive Models and Data Analytics in Software Engineering, PROMISE 2024, Porto de Galinhas, Brazil, 16 July 2024*, Weiwei Shang, Maxime Lamothe, and Zhiyuan Wan (Eds.). ACM, 56–65. <https://doi.org/10.1145/3663533.3664040>
- [3] Surya Mattu Julia Angwin, Jeff Larson and Lauren Kirchner. 2021. Machine Bias. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. Online.
- [4] Varya Monjezi, Ashish Kumar, Gang Tan, Ashutosh Trivedi, and Saeid Tizpaz-Niari. 2024. Causal Graph Fuzzing for Fair ML Software Development. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings, ICSE Companion 2024, Lisbon, Portugal, April 14-20, 2024*. ACM, 402–403. <https://doi.org/10.1145/3639478.3643530>
- [5] Varya Monjezi, Ashutosh Trivedi, Gang Tan, and Saeid Tizpaz-Niari. 2023. Information-theoretic testing and debugging of fairness defects in deep neural networks. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 1571–1582.
- [6] ProPublica. 2021. Compas Software Analysis. <https://github.com/propublica/compas-analysis>. Online.
- [7] Saeid Tizpaz-Niari, Ashish Kumar, Gang Tan, and Ashutosh Trivedi. 2022. Fairness-aware configuration of machine learning libraries. In *Proceedings of the 44th International Conference on Software Engineering*, 909–920.
- [8] Normen Yu, Luciana Carreon, Gang Tan, and Saeid Tizpaz-Niari. 2024. FairLay-ML: Intuitive Debugging of Fairness in Data-Driven Social-Critical Software. *CoRR abs/2407.01423* (2024). <https://doi.org/10.48550/ARXIV.2407.01423> arXiv:2407.01423
- [9] Normen Yu, Gang Tan, and Saeid Tizpaz-Niari. 2023. FairLay-ML: Intuitive Remedies for Unfairness in Data-Driven Social-Critical Algorithms. *CoRR abs/2307.05029* (2023). <https://doi.org/10.48550/ARXIV.2307.05029> arXiv:2307.05029