

Activity 6:

Introduction to the COMPAS Recidivism Algorithm

Lesson (4 hours)

Objectives

- Inform students of the harms of AI bias in the context of incarceration.
- Explain the historical bias in data.
- Students will learn about the Responsible AI Lifecycle.
- Understand why COMPAS is an example of AI Bias.

Resources

- [Slides for this activity](#). The file name for the slides is titled, “COMPAS.pdf”. The slides can be utilized to assist students with following along with the instructor.
- Here is the [blank Analyzing COMPAS assignment link](#) that should be given to students. Also, people can use the downloaded version by clicking on the “Blank_Deepnote_Analyzing_COMPAS” folder. Then click on “notebook.ipynb” to open the file on Visual Studio Code. Currently, we are working on developing a Pandas tutorial, so students feel prepared to complete the COMPAS activity.
- Here is the Answer key [link](#) for the Analyzing COMPAS assignment. Also, people can use the downloaded version by clicking on the “Answer_Key_for_Analyzing_COMPAS” folder. Then click on “notebook.ipynb” to open the file on Visual Studio Code.
- This is a [Cheat Sheet for Python basics](#) that helps students review the basics of Python and Pandas for data science. Also, people can use the downloaded version by clicking on the “cheat_sheet.pdf” file.

Preparation

Material	Description
Deepnote	Students can use an online notebook, such as Deepnote or Google Collab to complete the coding portion of this activity.

Visual Studio Code	If students don't want to use Deepnote to complete the coding portion of the project, students can download Visual Studio Code and open the "ipynb" files on Visual Studio Code.
Slides	The file names for the slides are titled, "COMPAS_slides.pdf"

Introduction

This lesson builds upon the previous activities because students will be required to understand how an AI algorithm, such as COMPAS can be bias. Students will act as data scientists for the company that created COMPAS, Northpointe. They will have to choose the input data for the AI algorithm called, COMPAS, which predicts whether someone is likely to commit a crime again. The activity will prepare students to complete the Intersectional data analysis of COMPAS on Deepnote because they will understand the results that they are trying to reproduce. The instructions correspond with each slide in the presentation.

Instructions

Activity 1: COMPAS Background [45 minutes]	
Instructions	
COMPAS Questionnaire[45 minutes]	<ul style="list-style-type: none"> ● Inform students that they will be tasked with deciding which data should be used to train COMPAS. ● Explain the background of COMPAS. ● Task 1 (20 minutes): Determine with your group the questions that should be used to decide whether someone needs to stay in jail by reviewing all 137 questions from the COMPAS Risk Assessment Questionnaire. Determine which questions should be given more weight in COMPAS because they are more indicative of someone committing a crime again. Here is the link to the COMPAS Risk Assessment Questions, which has the questions that

	<p>each person had to fill out before the COMPAS model made a prediction about their likelihood of committing another crime.</p> <ul style="list-style-type: none"> ● Task 2 (20 minutes): Each group shares the questions their group believes should be inputted into the COMPAS algorithm and the reason why they chose those answers. <ul style="list-style-type: none"> ○ Next, share the questions that contain information about a prior arrest, and therefore should not be used as input data for the COMPAS algorithm because there are racial disparities in arrests in the US [1]. These questions should not be the input data because they contain information about someone’s prior arrests: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11-19, 20-22, 23, 25, 26, 27, 28, 29, 30, 43, 44, 45, 46, 47, 48, 49, 50, 53, 54, 66, 133, 137 ○ Afterwards, the teacher will discuss how the company that made COMPAS, Northpointe, did not release which questions were used in COMPAS and which questions were given more weight in COMPAS. Inform students how this can be problematic, because it is a black box algorithm.
--	---

<p>Activity 1: Responsible AI Lifecycle [45 minutes]</p>	
<p>Instructions</p>	
<p>Responsible AI Lifecycle[40 minutes]</p>	<ul style="list-style-type: none"> ● Discuss how AI is not less biased than humans because data reflects human bias and we have missing data [2]. ● Review the Responsible AI Lifecycle, which is a series of questions that data scientists should ask themselves while they are making algorithms.

	<ul style="list-style-type: none"> ● Introduce the “Intersectional data analysis of COMPAS” activity. During this activity, students will use Python and Pandas(a data analysis library) to analyze COMPAS. Additionally, students will learn about three different fairness metrics: accuracy, demographic parity, and false negative/false positive parity. We will use an interactive data science notebook called “Deepnote”, so students can quickly run each line of code. Their task is to determine the fairness metric that should be used to determine whether COMPAS is fair. Moreover, they will regenerate similar results that investigative journalists from Propublica discovered about COMPAS. For example, Propublica found that “Black people are almost twice as likely as White people to be labeled a higher risk but not actually re-offend ... It makes the opposite mistake among White people: They are much more likely than Black people to be labeled lower risk but go on to commit other crimes” (Propublica). The “Intersectional data analysis of COMPAS” activity will take about 2 hours. <ul style="list-style-type: none"> ○ Here is the blank assignment link that should be given to students. ○ Here is the answer key link
Reflection[5 minutes]	<ul style="list-style-type: none"> ● Share your experience doing the activity: <ul style="list-style-type: none"> ○ What did you learn? ○ What was challenging?

References

[1] Alexander, Michelle. *The New Jim Crow: Mass Incarceration in the Age of Colorblindness*. New York, The New Press, 2010.

[2] Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity.

[3] Jeff Larson, Julia Angwin. “How We Analyzed the Compas Recidivism Algorithm.”
ProPublica, 23 May 2016,
<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.