



Racial Bias in Healthcare Data Solutions

Perpetuating Disparities in Medical Treatment Nationwide

Themes:

Algorithmic Bias & Fairness
Social & Economic Impact

Prerequisites:

- None for the Case Study section.

Owner:

[Center for AI and Data Ethics](#) at University of San Francisco

Author(s):

Hadley Dixon and Robert Clements

License:

Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International
[CC BY-NC-SA](#)

Citation:

Dixon, Hadley and Clements, Robert. (2024). Racial Bias in Healthcare: Perpetuating Disparities in Medical Treatment Nationwide.

Objective:

The purpose of this case study is to use a role-based approach to examine the different perspectives and incentives behind both the development of a data-driven

healthcare application that was discovered to be unfair to people of color, and the corresponding reaction of the company.

Instructions:

1. Students should form a group of 3 or 4.
2. Non-technical audiences: Each person in the group should be assigned one or two roles, depending on the size of the group. Once roles are assigned:
 - a. Each group member should read the scenario and the description of each role.
 - b. The student in the Moderator role will moderate the conversation, while also playing their own role during the discussion.
 - c. The Moderator should ask each student to introduce their role.
 - d. The Moderator should kick off the discussion using the list of discussion questions as a guide. The discussion should be allowed to go in different directions organically, and the listed questions can be used if there is a need to get back on track.
3. Technical audiences: each person in the group should be assigned one or two roles, depending on the size of the group, except for the Moderator (CEO of Insurance Client) role. Instead, the Product Manager should now act as the Moderator. If there are more than three people in the group, assign each extra group member to be a Data Scientist reporting to the Data Science Manager.
 - a. Follow (a)-(d) from Step 2 above, but now the discussion should be much more technical in nature.
4. Technical audiences: complete the Colab notebook.

Case Study:

In recent years, a diverse set of healthcare entities, such as insurance firms, technology providers, pharmaceutical giants, biotech innovators, and hospital networks, have increasingly adopted artificial intelligence systems and machine learning algorithms for different parts of their business. This integration aims to enhance diagnostic precision, build and improve predictive analytics, and optimize resource allocation to enhance the quality of healthcare services, improve patient outcomes, streamline processes, identify wasteful spending, among other priorities. While the use of AI in patient care brings significant benefits, it is crucial to address ethical considerations and perform rigorous examination of existing models. “Bias” in healthcare models can be especially harmful and contribute to existing patient inequities. We put the word “bias” in quotes because bias can have multiple meanings. Moving forward, we define bias to be a model that exhibits unfairness to groups of people according to specific fairness metrics.

In the following scenario, Optumize Technologies, a leading healthcare technology solutions provider, partners with an insurance client to enhance patient care and optimize resource utilization for their members - a diverse patient population with varying medical needs. The insurance client faces a significant challenge in efficiently identifying individuals who are at a higher risk and in greater need of medical

assistance. The rising prevalence of chronic conditions among policyholders increases associated healthcare costs, necessitating an innovative approach to risk stratification. The client aims to deploy an advanced strategy that analyzes policyholder data to flag individuals in need of additional medical support to better coordinate preventative care and reduce future costs.

A team of data scientists within Optumize Technologies is tasked with designing a predictive model that will use policyholder data provided by the insurance client to identify patients that may need more medical interventions in the future. The team is excited to contribute to healthcare advancements and improve patient outcomes through the use of machine learning (ML) algorithms. The model should function by flagging patients in need of additional medical support as “high-risk”. To determine patients’ risk level, the model analyzes comprehensive medical histories, including visits and utilization, medical procedures, and costs.

One issue that the data science team uncovers is in the company’s definition of “high risk” because a suitable measure of risk does not already exist. The team decides to use a patient’s historical health care costs (extracted from insurance claim data) as a proxy for risk level and they proceed to develop a model that predicts future medical costs - the higher the predicted costs, the more risky the patient is predicted to be.

Following development and deployment of the model, which is sold as a product under the name HealthOptumizer, it has become one of the most widely used data-driven healthcare products across the country in a matter of months, due to the expansive insurance clientele of Optumize Technologies. An estimated 150 million patients are served each year, across the private and public sector. Over the past three years, continuous monitoring and updates have been made to adapt the underlying model in HealthOptumizer to evolving data patterns.

An independent research team conducts a bias and fairness study on HealthOptumizer. Their findings reveal significant racial bias in the model, which impacts the identification of high-risk patients, favoring white patients over black patients. The research team investigates the root cause of this bias and assesses its implications on healthcare disparities.

The data the research team was able to obtain, comprising 6,000+ black individuals and 44,000 white individuals, showed that black patients had 24% more chronic health conditions than their counterparts. Despite this, the study revealed a consistent pattern where the algorithm disproportionately categorizes white patients as “high-risk”, even when black patients exhibit more severe health conditions. Upon deeper investigation, the team identified that the algorithm primarily relies on predicting future expenditures rather than focusing on the severity of the illness. White patients historically receive more medical care, leading to the skewed risk prediction.

Moderator and CEO of Insurance Client

The CEO of the insurance client has called together all of the parties involved in order to discuss the problems with HealthOptumizer and the available solutions. The CEO is not a technical expert in ML, and is at the meeting primarily to hear about what can immediately be done to fix the current issues, and how to prevent them from

happening again. They are not interested in excuses, or assigning blame. They just want to move forward as quickly as possible with a solution.

Data Science Manager at Optumize Technologies

The Data Science Manager has worked in the data science field, at several companies and in several roles, both hands-on and managerial, for over a decade. They manage the team of data scientists that developed this model, and take full responsibility for the damage it has caused to their company's and their team's reputation. They recognize that their team failed to assess the model completely. It was not common practice to test for algorithmic fairness for any of their models in the past, and it is also not something that they have done in previous roles at other companies. While they are proud of the work their team does, they realize that sometimes the speed at which they are expected to deliver results can impact the quality of the final product in many ways.

Research Team Member

The Research Team Member, not employed at either Optumize Technologies or the insurance client, helped uncover the problems with the model, and has been invited to the meeting in order to provide any guidance they can. They work as part of an academic team investigating algorithmic bias and fairness in ML models, and care deeply about the impacts that commercial products can have on society, especially people who are at a disadvantage. They also approach new technology suspiciously, and feel that for-profit companies should be more transparent with their customers about what they are doing with their customer's data. Note that their analysis is not under question here, we can assume that the problems with the model are very real.

Product Manager at Optumize Technologies

The Product Manager was managing the development and release of HealthOptumizer, along with a portfolio of other AI-based products. They are very proud of their work, and their ability to meet the demands and expectations of the customer (the insurance client) and their own manager. This product in particular was extremely valuable to the company, and there was considerable pressure from the customer to release it on time. They are also aware that there is another team within the company, with more face time with senior leaders, that has been promising to build a similar product. Getting this product to market sooner brought considerable positive attention to a lesser known data science team.

Discussion Questions:

1. Each person at the meeting, sequentially, should provide a recommendation for dealing with HealthOptimizer right now, and what should be done about it in the near future?
 2. Each person should provide a recommendation for avoiding the same issue with a future product at Optimize Technologies.
 3. Each person should carefully assess their own role, if any, in the failure of HealthOptimizer.
-

References:

Ziad Obermeyer et al., Dissecting racial bias in an algorithm used to manage the health of populations. Science 366, 447-453 (2019). DOI:[10.1126/science.aax2342](https://doi.org/10.1126/science.aax2342)

