




Case3 Transcript



#Slide1

This third case for the algorithmic bias module is on predicting acute kidney injury.

#Slide2

In this scenario, consider an algorithm that was developed by a major technology company to predict the onset of acute kidney injury in patients within the next 48 hours.

#Slide3

The algorithm was developed using a set of EHR data from the US Department of Veterans Affairs with 703,782 adult patients.

#Slide4

AKI is often thought to be preventable with early treatment, but current algorithms for detecting AKI depend on changes in serum creatinine which lags behind the actual decline in renal function leading to delays in treatment. This algorithm was able to predict 55% of inpatient AKI events of any severity within a window of up to 48 hours in advance. It also provided correct early predictions in 84% of episodes in which inpatient or outpatient dialysis would be required within 30 days of the onset of AKI and 90% of cases in which outpatient dialysis would be scheduled within 90 days.

#Slide5

The authors admit that the model is not representative of the population with only 6% of the data comprised of female patients, for whom the model's performance was lower.

#Slide6

Despite this, the company reported on the model in a blog post which did not address this limitation at all.

#Slide7

Based on the information provided in the case, please take a moment to answer at least one the following questions linked below the case video:

How would you deal with advertising from companies promoting their AI products?

What are the pros and cons of using VA data to train AI algorithms?