

## Translational Tutorial: Understanding the context and consequences of pre-trial detention

Elizabeth Bender (*Staff Attorney, The Legal Aid Society of NYC / Decarceration Project*)

Kristian Lum (*Lead Statistician, Human Rights Data Analysis Group*)

Terrence Wilkerson (*entrepreneur*)

The 2016 ProPublica article, [Machine Bias](#), brought the uses of predictive modeling in the criminal justice system to the forefront of popular attention and centered the discussion around racial bias in the software used to make the predictions. This stimulated a lively debate about what it means for predictions to be fair with respect to a protected class, such as race. The article focused on the use of COMPAS, a software used to make recommendations to judges regarding whether a defendant should be granted pre-trial release. The simultaneous release of the data on which ProPublica's analysis was founded allowed machine learning and other quantitative researchers to engage in this debate, putting forward a variety of (often competing) notions of fairness and accompanying computational methods to achieve the proposed definitions. In fact, recapitulations of the fairness of the ProPublica compas dataset have been a popular theme in past FAT/ML conferences, comprising 5 of 13 talks and papers in 2016.

In order to address the fairness of COMPAS's predictions in a framework familiar to machine learning researchers, the details surrounding the system generating the data and how the recommendations will affect the individuals' lives are somewhat abstracted away. Individuals and their stories become counts in a confusion matrix. But, without a deep understanding of the context in which the predictions will be used, we believe it is difficult, if not impossible, to make meaningful determinations of the "fairness" of the predictions and derived recommendations. How can one declare recommendations for pre-trial detention "fair" (regardless of equality of some metric among racial classifications) without fully understanding the consequences of pre-trial detention on the lives of the people about whom the recommendations are made?

In this tutorial, we hope to paint a richer picture of context in which pre-trial detention decisions are determined and how recommendations for pre-trial detention might affect individuals' lives. We will begin with a brief overview of risk assessment tools and their varied uses. We will also cover the process from arrest to disposition from a legal and procedural perspective. Finally, Terrence Wilkerson-- a person twice falsely accused of armed robbery who has spent a total of two years in pre-trial detention-- will share some of his experiences with the criminal justice system and explain how bail and the resultant pre-trial detention affected his case outcome and life.

This tutorial is designed for quantitative researchers who work with the ProPublica COMPAS dataset or other criminal risk assessment datasets and would like a more in-depth understanding of the legal processes that generate the datasets and the experiences and stories that make up the datasets but aren't told by the data itself.