Using Data Science to Achieve Fair and Equitable Outcomes

Rayid Ghani

Carnegie Mellon University

MACHINE LEARNING DEPARTMENT

Heinz College
INFORMATION SYSTEMS • PUBLIC POLICY • MANAGEMENT

Reducing jail recidivism with proactive mental health interventions (Johnson County, KS)

Reducing Incarceration through Prioritized Interventions. Bauman et. Al. ACM COMPASS 2018

11 MILLION

people move through 3,100 Jails

22 BILLION

in cost

64%

suffer from mental illness

68%

have a substance abuse disorder

44%

suffer from chronic health problems

Children in at least 4MM US households are exposed to high levels of lead

Impaired *Attention*

Hearing Loss

Lower IQ

Lack of Motor Skills

Learning Disability

Memory Problems

Increasing Retention in Care for HIV+ Patients

Predictive Analytics for Retention in Care in an Urban HIV Clinic. Ramachandran et al. Nature Scientific Reports 2020

HEALTHY CHICAGO
CHICAGO DEPARTMENT OF PUBLIC HEALTH

CCHE
The Chicago Center for HIV Elimination

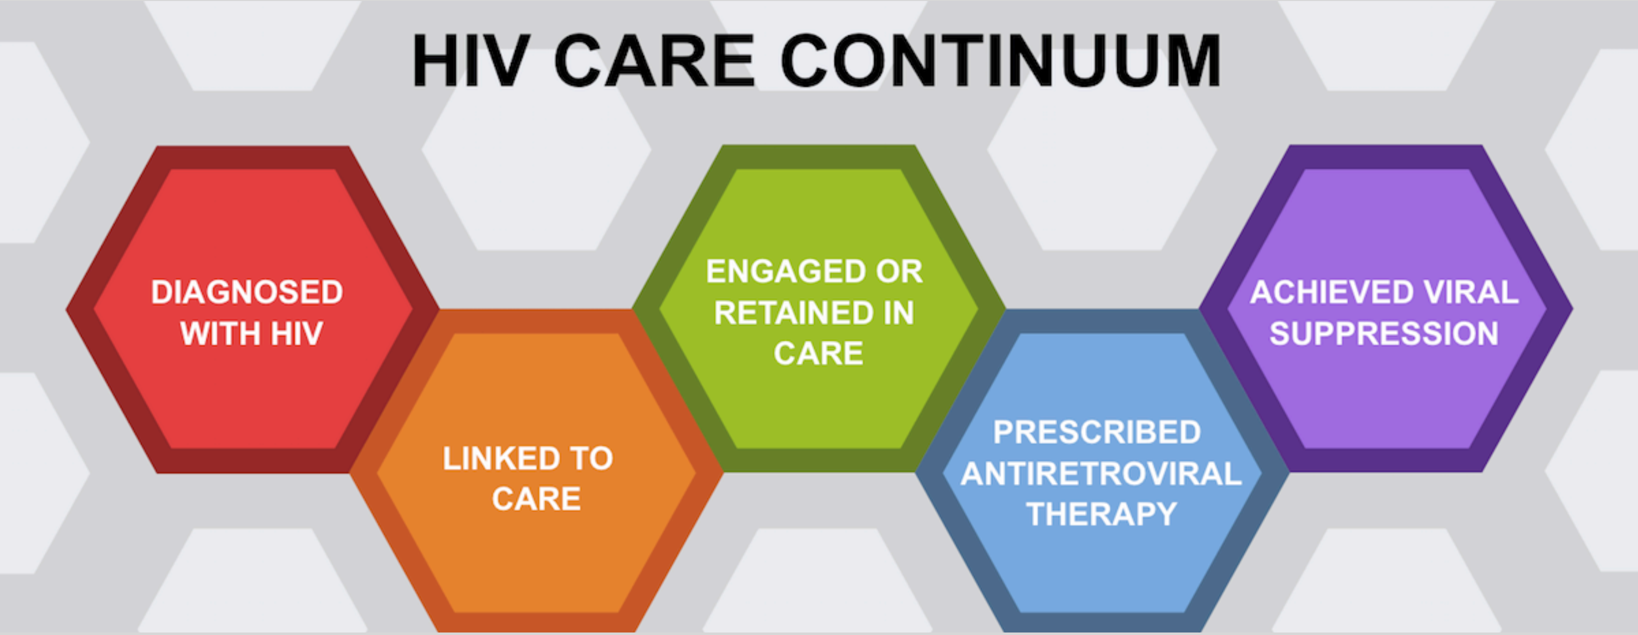

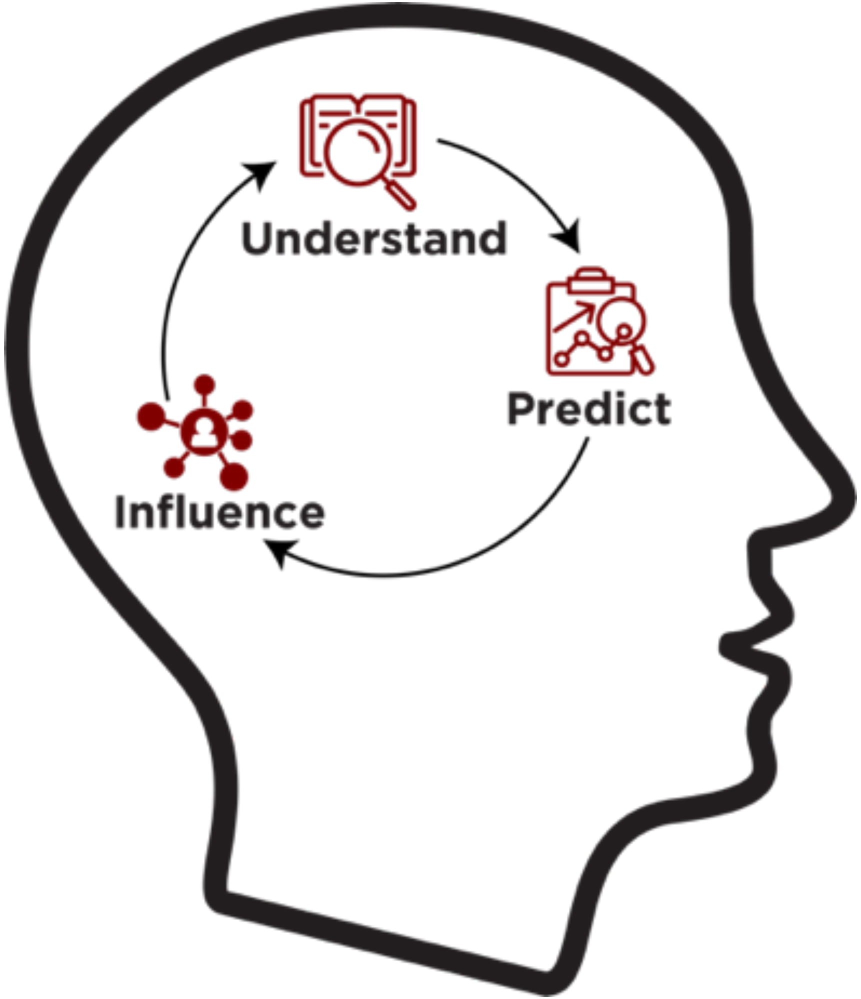
Understand
Predict
Influence

How do we develop Human-ML collaborative systems to **help make decisions** that lead to **fair** and **equitable outcomes**?

Human-AI Collaborative Systems

Allocation of Limited Resources

Balancing goals of equity, efficiency, and effectiveness

The focus is not just be on making the ML model fair but rather on **making the overall system and outcomes fair**

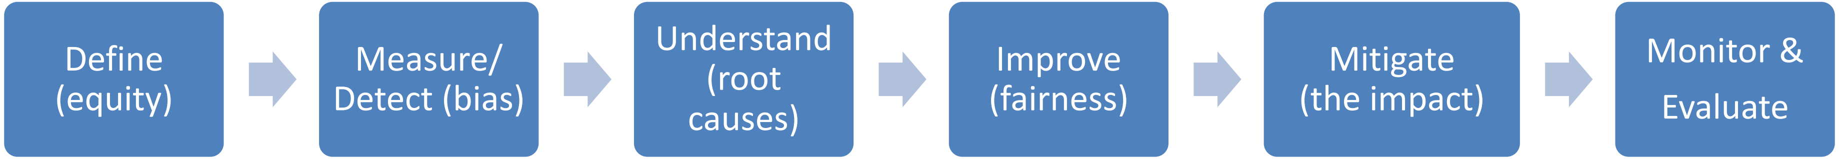

Bias (in outcomes) can come from any of these four components

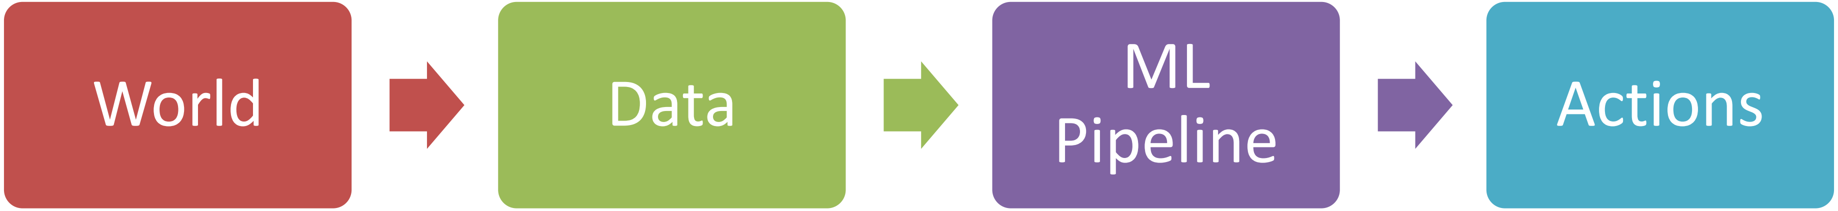

Sample Bias
Measurement Bias
Label Bias

System Developers
Complexity or flaws
Design Choices

Many Bias Measures: How do we select what we care about?

- Statistical/Demographic Parity
- Impact Parity
- False Discovery Rate Parity
- False Omission Rate Parity
- False Positive Rate Parity
- False Negative Rate Parity
- …

FAIRNESS TREE

Joint Work with Kit Rodolfa and Pedro Saleiro

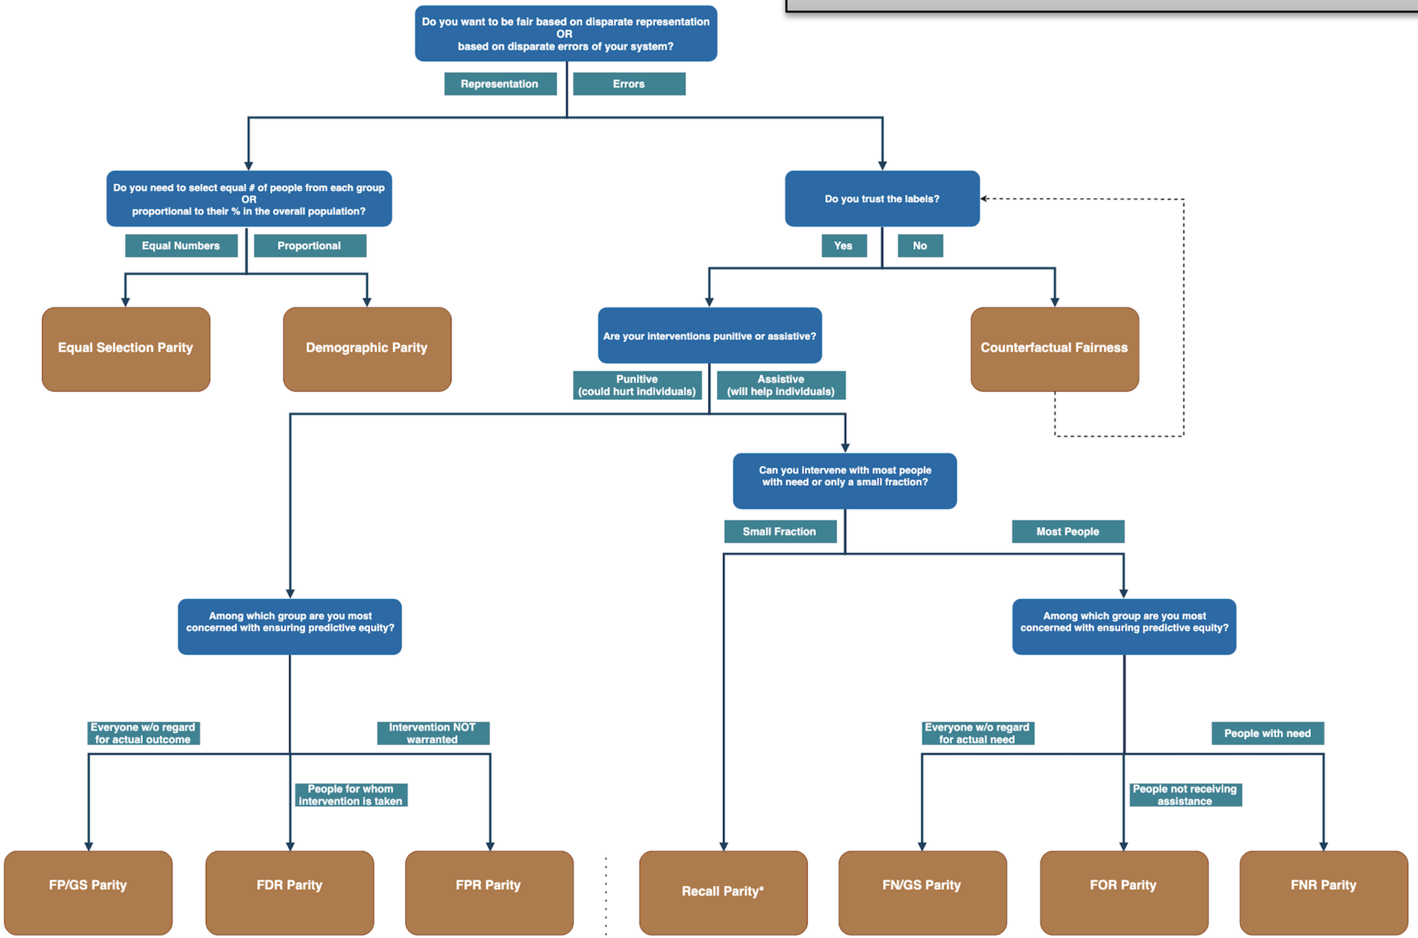

Zoomed in Version

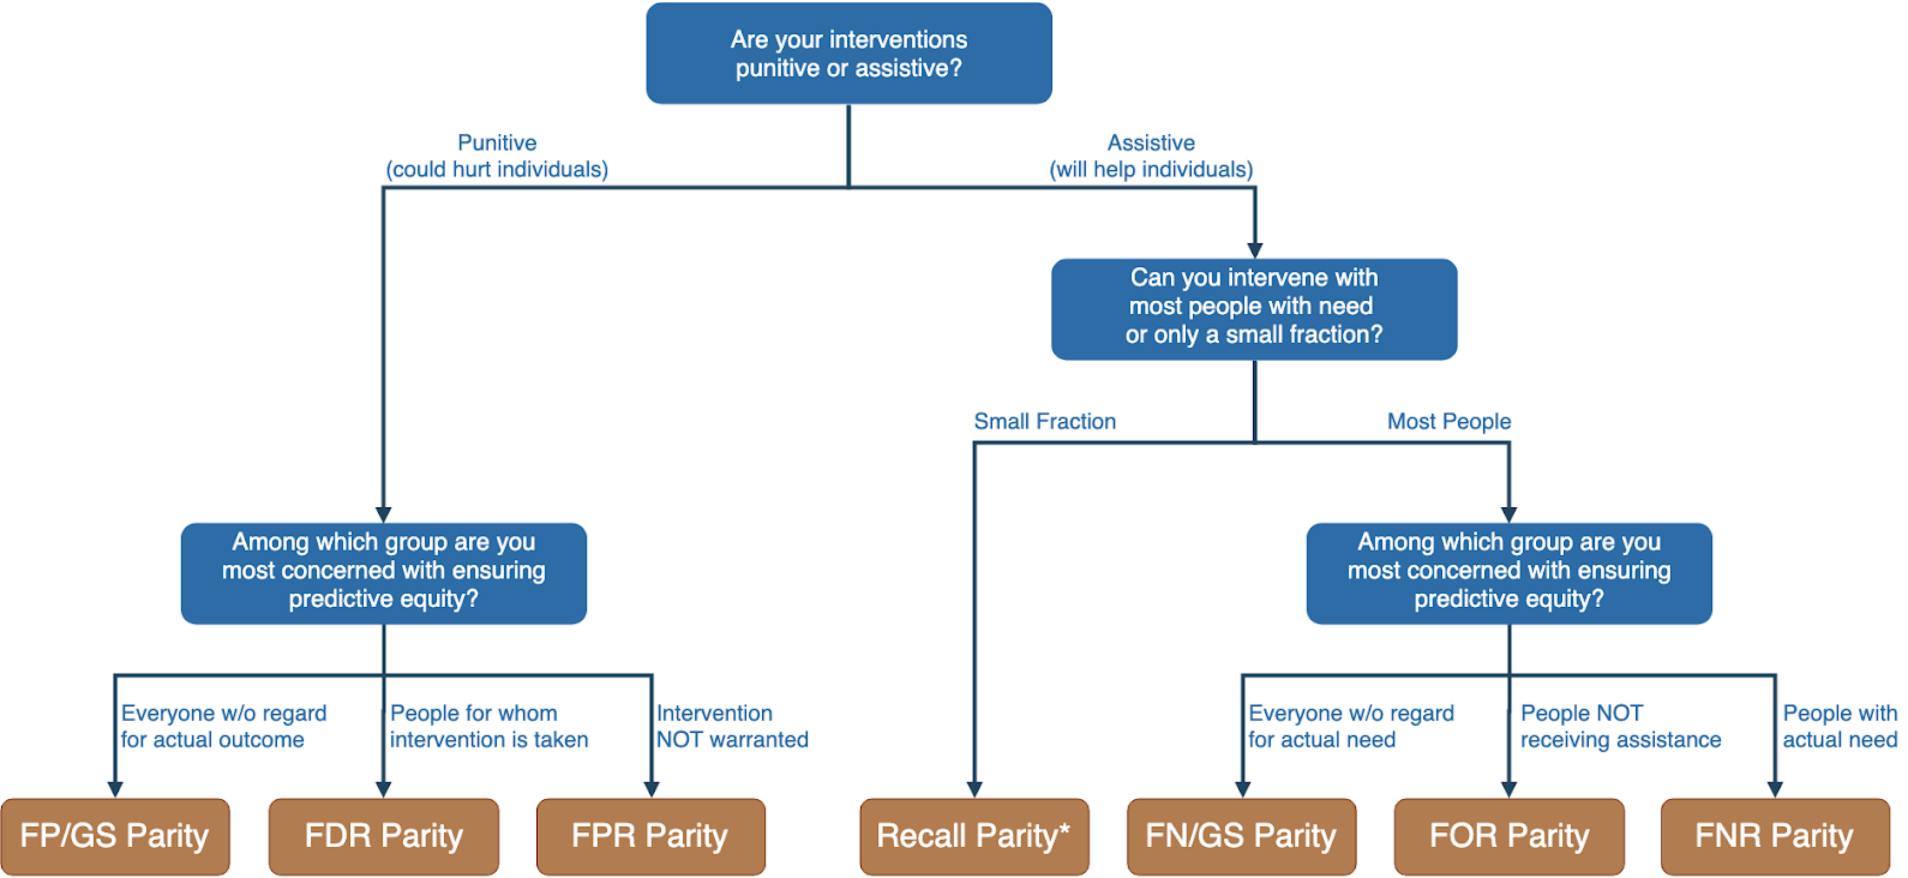

Aequitas

Open Source Bias & Fairness Audit Tool

Bias and Fairness Audit Report

Generated by Aequitas for [Large US City] Criminal Justice Project
January 29, 2018

Project Goal: Identify individuals likely to get booked/charged by police in the near future
Performance Metric: Accuracy (Precision) in the top 150 identified individuals
Bias Metrics Considered: Demographic Disparity, Impact Disparity, FPR Disparity, FNR Disparity, FOR Disparity, FDR Disparity
Reference Groups: Race/Ethnicity – White, Gender: Male, Age: None

Model Audited: #841 (Random Forest) **Model Performance**: 73%

⚠ Aequitas has found that Model 841 is **BIASED**. The Bias is in the following attributes:

Race = Black is **biased** in Demographic Disparity (6X), Impact Disparity 1.8X) , FPR Disparity (5X), FOR Disparity (1.5X) , FDR Disparity (1.7X)

46% (66) of the selected group (n=150), while only making up 24% of the total population.
FDR (30%) is 1.7X higher than Reference FDR (18%).
FOR (6%) is 1.5X higher than Reference FOR (4%).
FPR (0.02%) is 5X higher than Reference FPR (0.004%)


0%
100%
Performance
Metric

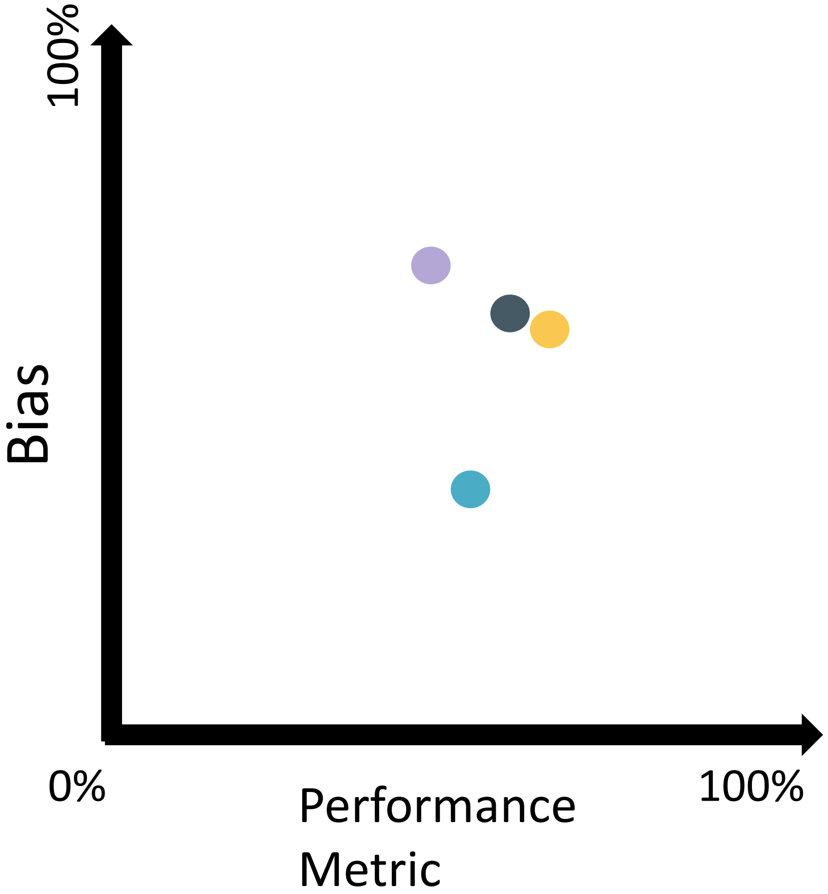
100%
Bias
0%
Performance Metric
100%

Case Study: Prioritizing Early Diabetes Screening

Data from 1M patients from 2006 to 2018

Outcome: Type 2 Diabetes diagnoses in the 3 years period after a provider visit.

Case Study: Prioritizing Early Diabetes Screening

Performance Metric: Recall/Sensitivity@k (k= # of patients screened based on resources)

Group Metric: False Omission Rate

Protected Attributes: Age, Ethnicity and Sex

Test Set:

- 205,485 patients that visit a doctor during 2014
- 3.4% prevalence (7,154 diabetes diagnoses in 3 years interval after visit)

Current Practice (at provider)

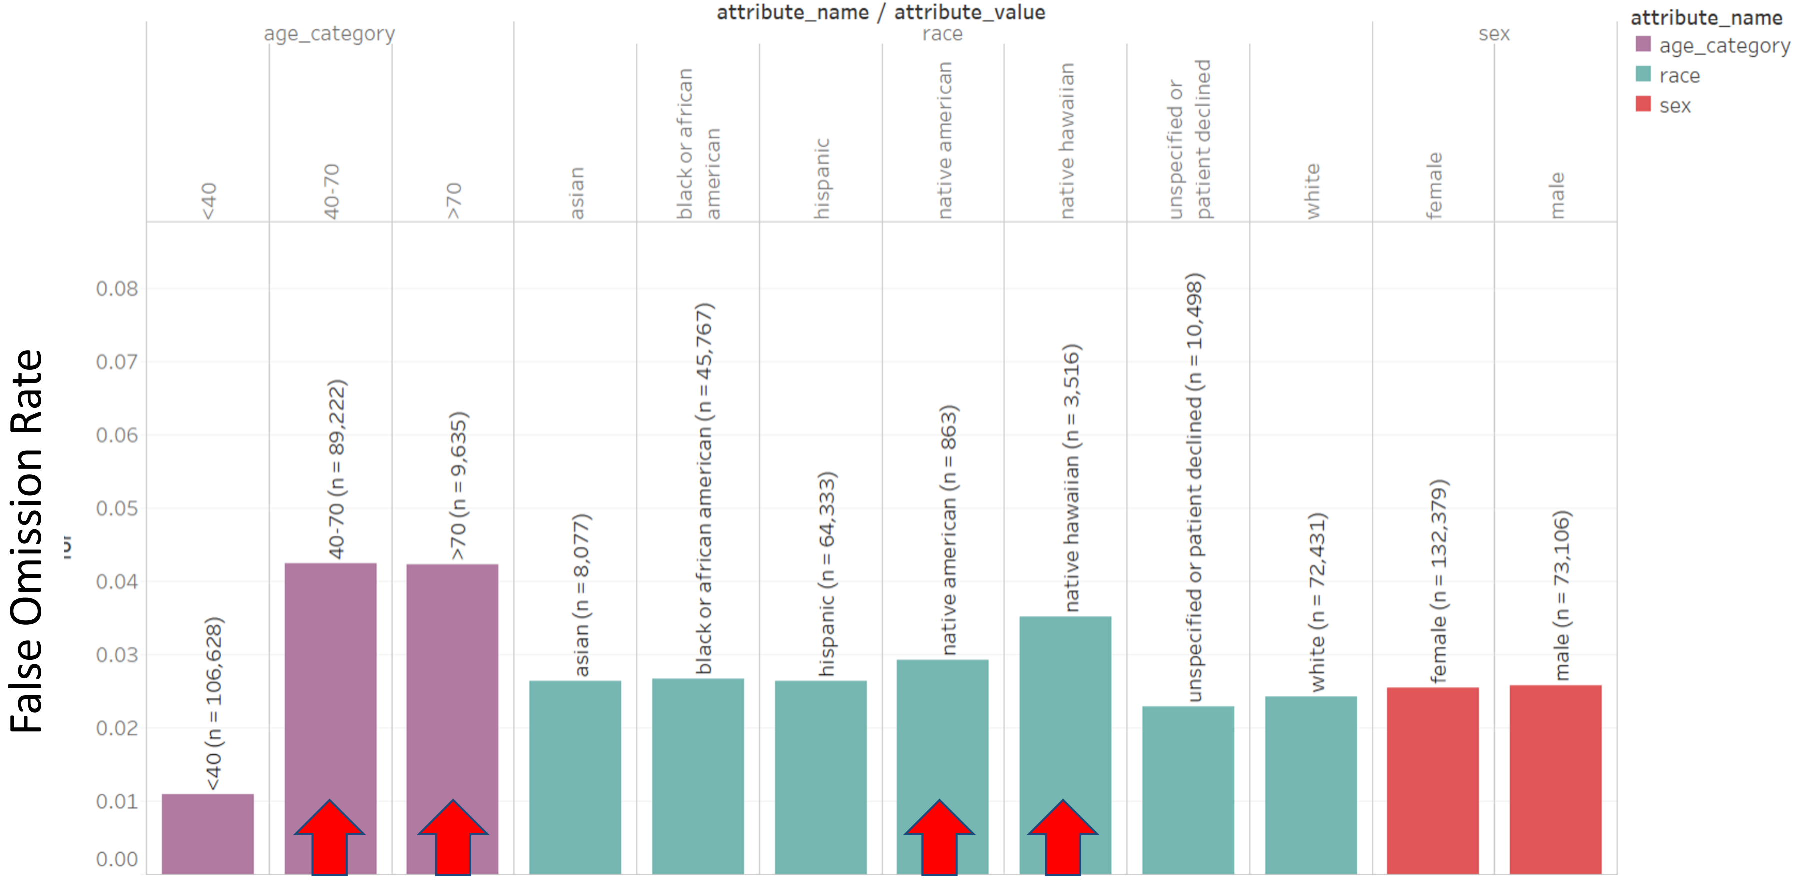

Current Practice vs "Best Overall" Model

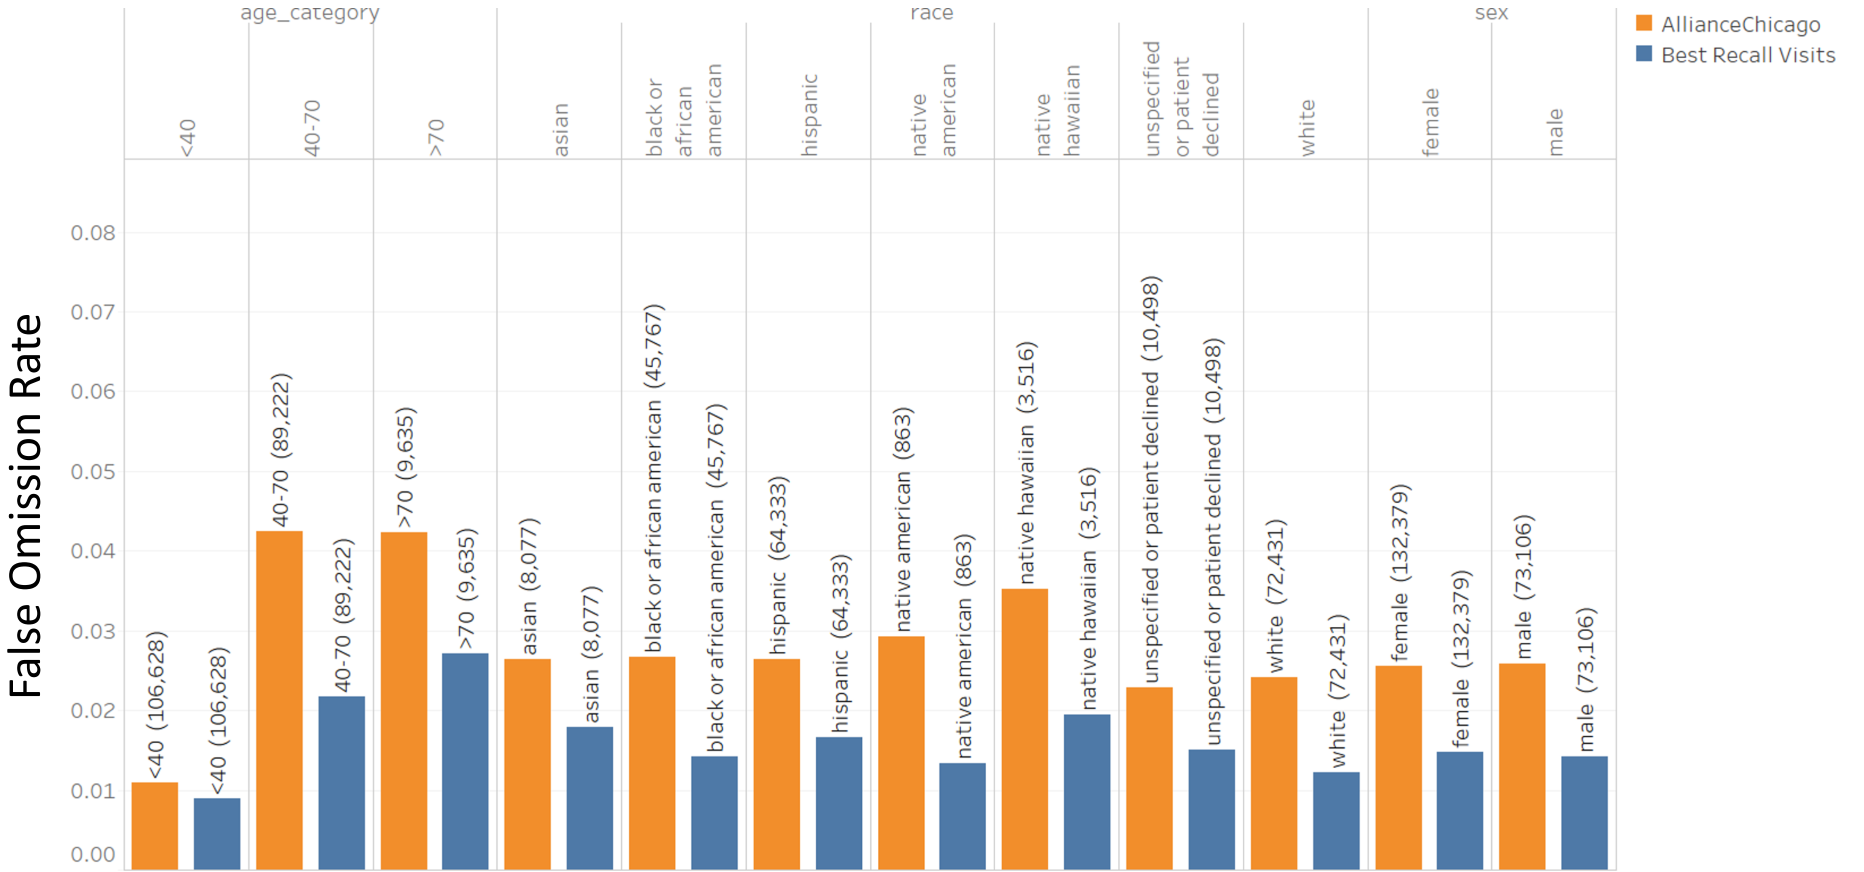

Current Practice vs "Best Overall" vs "Best Hispanic" Model

Current Practice vs “Best Overall” vs “Race Fair” Model

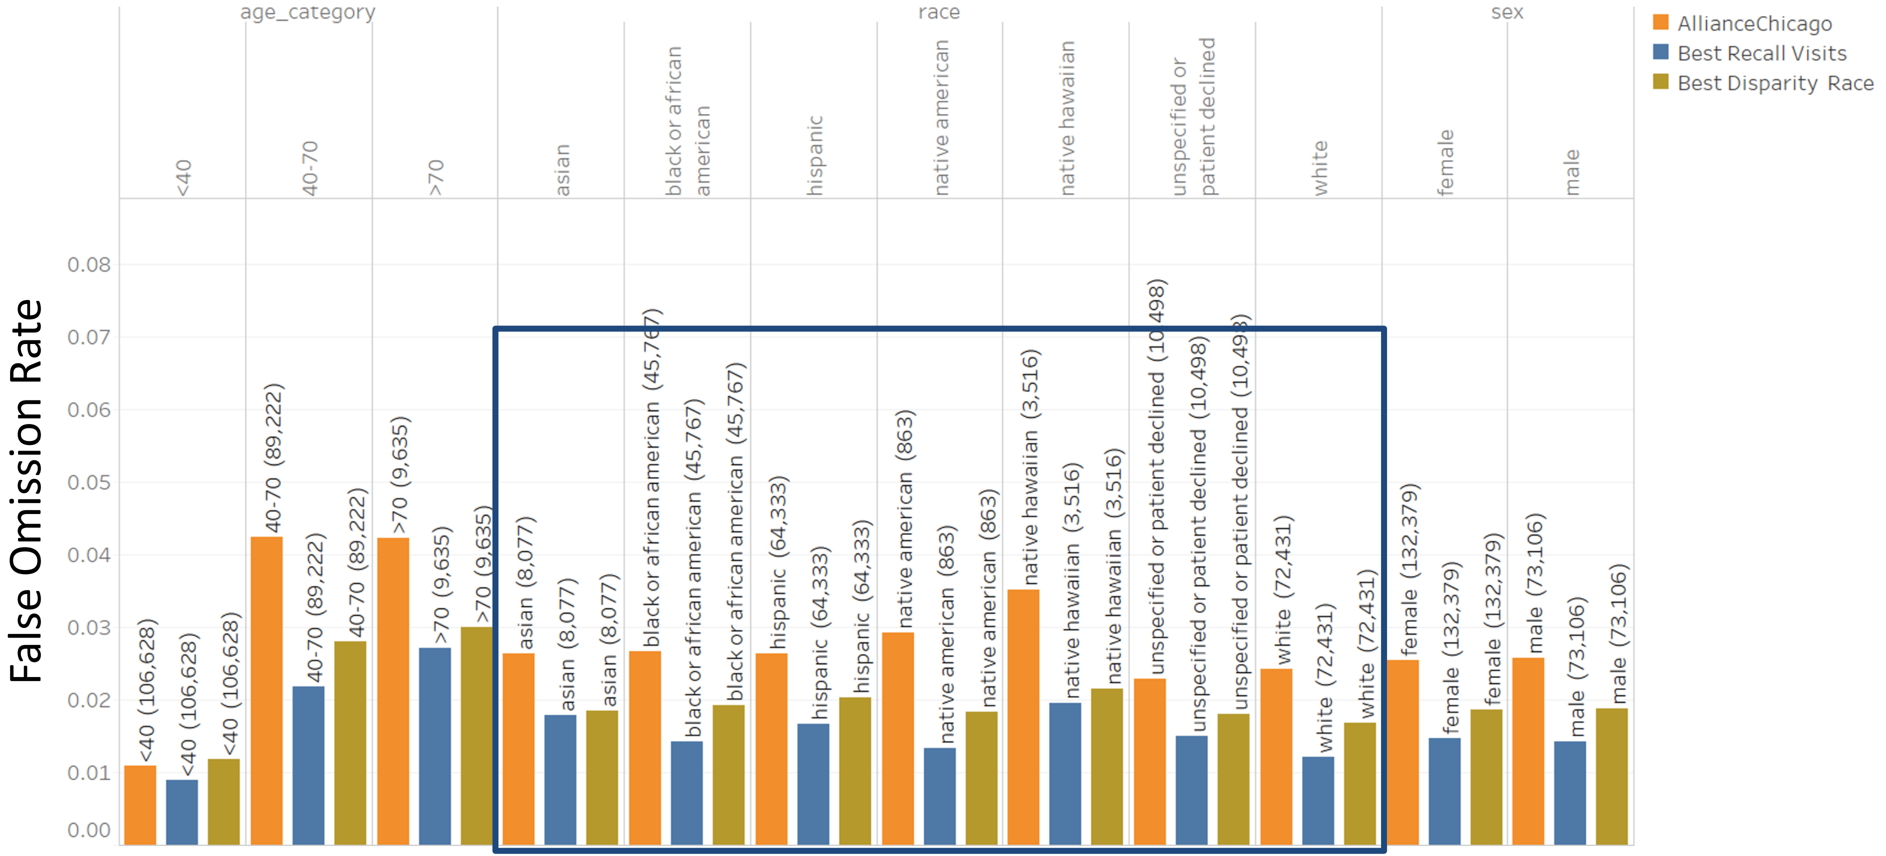

Case Study: Reducing Misdemeanor Recidivism through Diversion and Social Service Programs

Predictive Fairness to Reduce Misdemeanor Recidivism Through Social Service Interventions. Rodolfa et al. ACM FAT* 2020

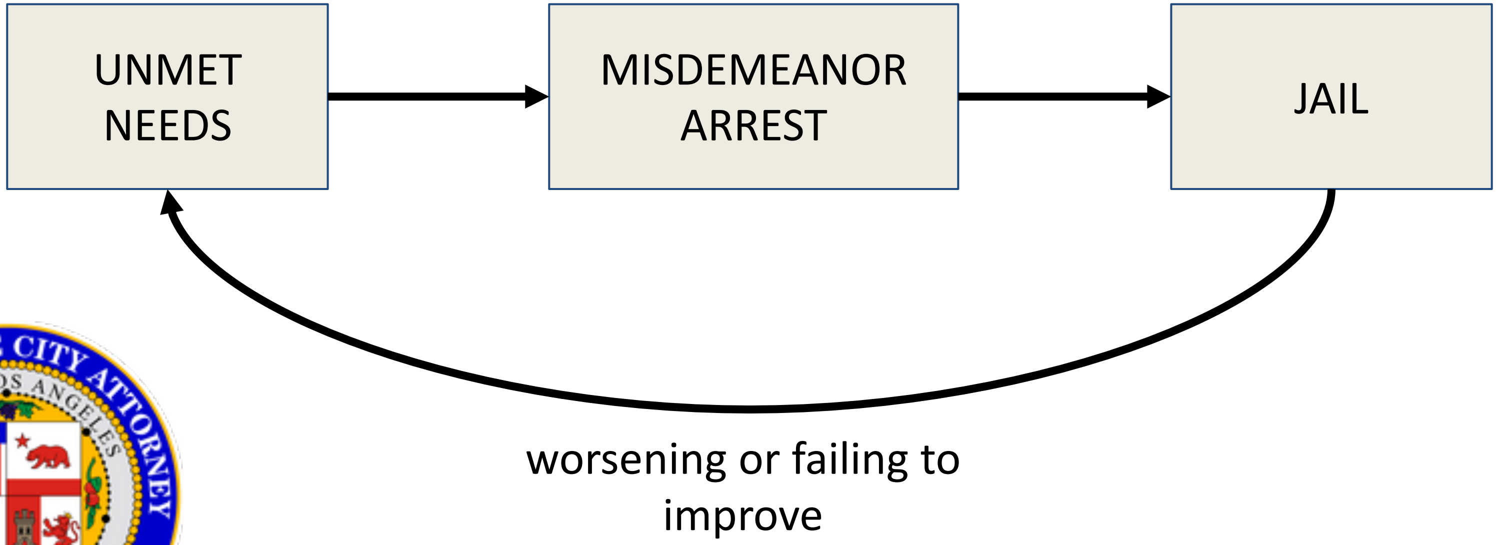

Case Study: Breaking the Cycle

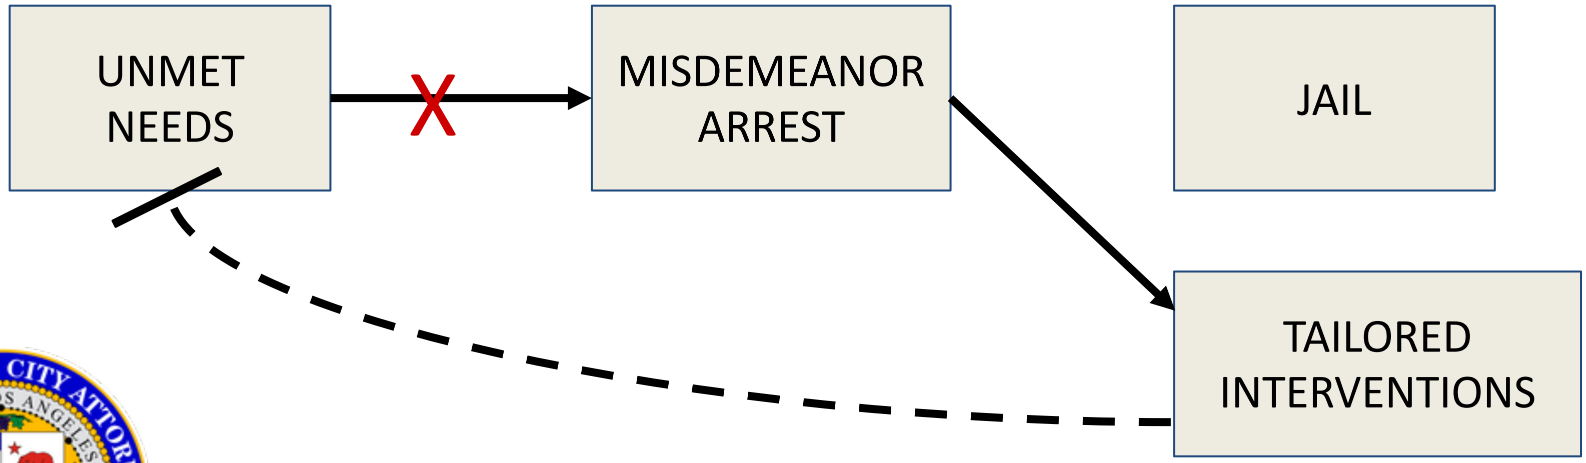

Policy Menu

Designing for Efficiency

72.7% Efficient

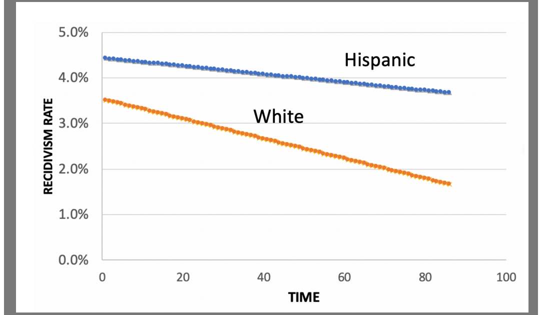

Equality

Additional Cost: 2%

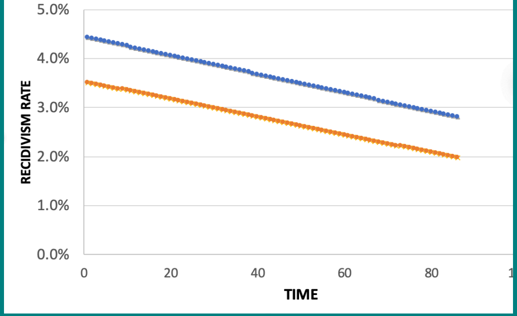

Equity

Additional Cost: 2%

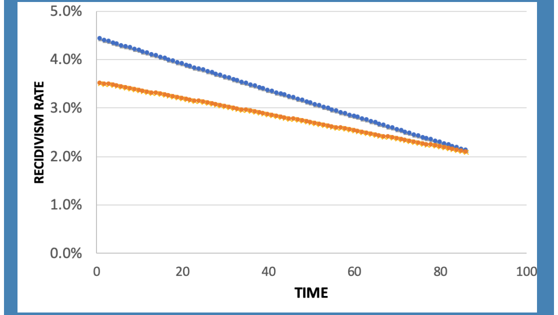

Rayid Ghani

Carnegie Mellon University

MACHINE LEARNING DEPARTMENT

Heinz College
INFORMATION SYSTEMS • PUBLIC POLICY • MANAGEMENT

rayid@cmu.edu

code at github.com/dssg