Statistical Methods for Algorithmic Fairness in Risk Adjustment

Sherri Rose, Ph.D.

Associate Professor of Health Policy
Co-Director, Health Policy Data Science Lab

Stanford University

drsherrirose.org
@sherrirose

November 17, 2021

H
ECONOMICS
A
POLICY
OUTCOMES
H

“ Learning two fields takes, surprisingly, twice as long as learning one. But it’s worth the investment because you get to solve real problems for the first time. ”

Barbara Engelhardt | Princeton

“ In both private enterprise and the public sector, research must be reflective of the society we’re serving. ”

Rediet Abebe | Harvard & UC Berkeley

“ …behind every data point there is a human story, there is a family, and there is suffering. ”

Nick Jewell | LSHTM & UC Berkeley

FAIRNESS

Who decides the research question?

Who is in the target population?

What do the data reflect?

How will the algorithm be assessed?

Who decides the research question?

Who is in the target population?

What do the data reflect?

How will the algorithm be assessed?

Justice: benefits, risks, costs, and resources are equitably distributed

Ethical Pipeline

Problem Selection

Disparities in funding and problem selection priorities are an ethical violation of principles of justice.

Ethical Pipeline

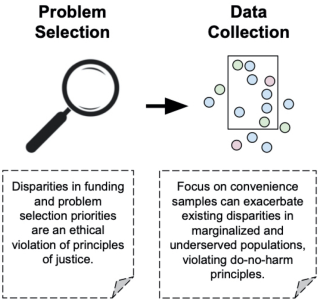

Ethical Pipeline

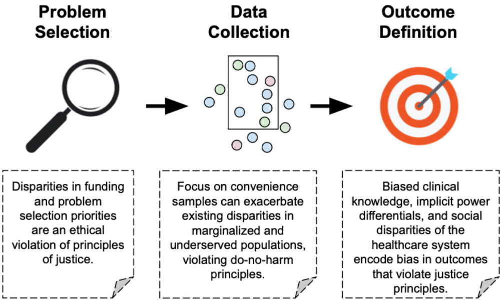

Ethical Pipeline

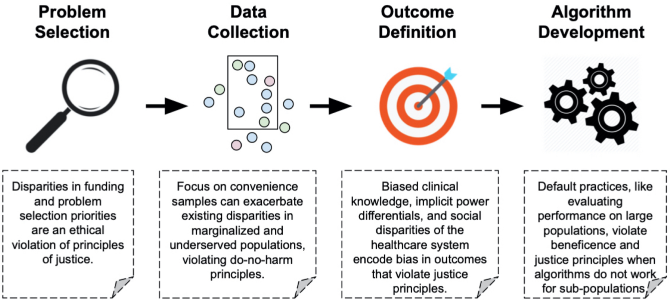

Ethical Pipeline

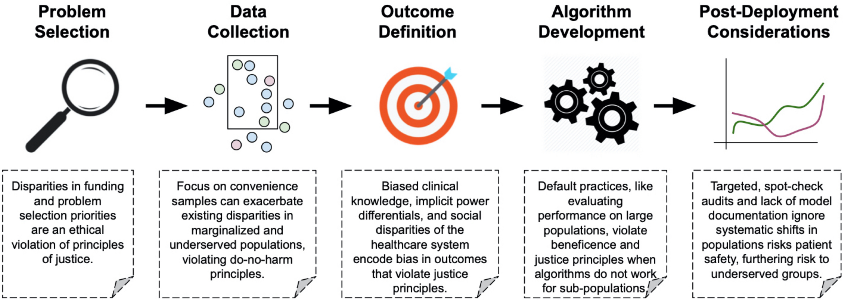

Ethical Pipeline

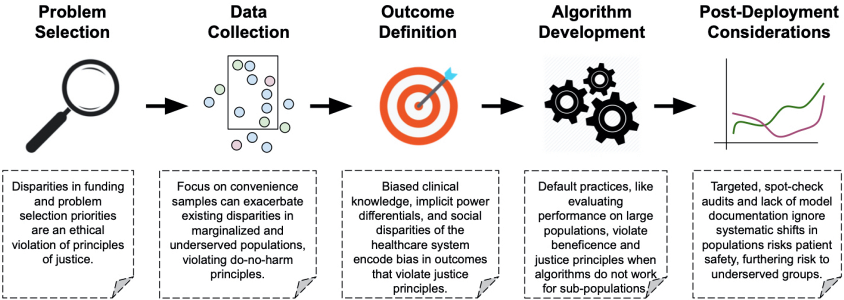

Irene Chen
PhD Student
MIT

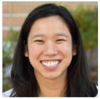

Ethical Machine Learning in Healthcare

Annual Review of Biomedical Data Science
Vol. 4:123-144 (Volume publication date July 2021)

Irene Y. Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

$$Y = \theta X$$

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Spending outcome

$$Y = \theta X$$

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Spending outcome

$$Y = \theta X$$

Input vector

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Spending outcome

$$Y = \theta X$$

Coefficient vector

Input vector

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Spending outcome

$$Y = \theta X$$

Coefficient vector

Input vector

Plan Payment Risk Adjustment

Over 50 million people in the United States currently enrolled in an insurance program that uses risk adjustment

- Redistribute funds based on health
- Encourage competition based on efficiency and quality
- Massive financial implications

Spending outcome

$$Y = \theta X$$

Coefficient vector

Input vector

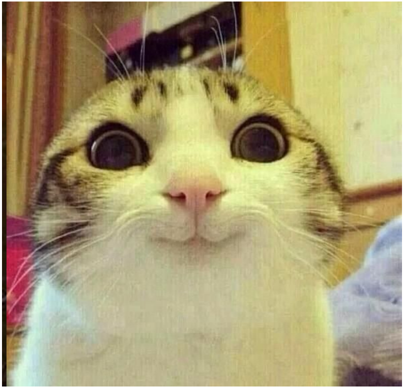

Variable Selection and Upcoding

Reduced set of 10 variables 92% as efficient

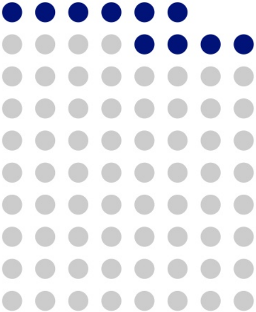

A Machine Learning Framework for
Plan Payment Risk Adjustment

Sherri Rose

Variable Selection and Upcoding

~~Reduced set of 10 variables 92% as efficient~~

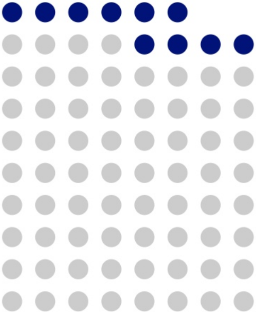

"…results for the risk adjustment algorithms that considered a limited subset of variables…performed consistently worse across all benchmarks."

Sample Selection for Medicare Risk Adjustment Due to Systematically Missing Data

Savannah L. Bergquist, Thomas G. McGuire, Timothy J. Layton, and Sherri Rose

A Machine Learning Framework for Plan Payment Risk Adjustment

Sherri Rose

Improving Mental Health Care, 1950-2000

Changes in financing and organization of mental health care, not new treatment technologies, made the difference

"Improvements ... evolved through ... more money, greater consumer choice, and the increased competition among ... providers that these forces unleashed"

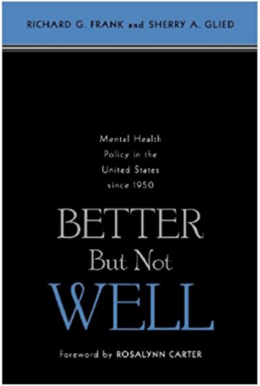

Mental Health and Substance Use Disorders (MHSUD)

Risk adjustment in the Marketplaces
recognizes only 20% of enrollees with MHSUD

Individuals with MHSUD can be **systematically discriminated** against

By Ellen Montz, Tim Layton, Alisa B. Busch, Randall P. Ellis, Sherri Rose, and Thomas G. McGuire

Risk-Adjustment Simulation: Plans May Have Incentives To Distort Mental Health And Substance Use Coverage

Fairness Methodology

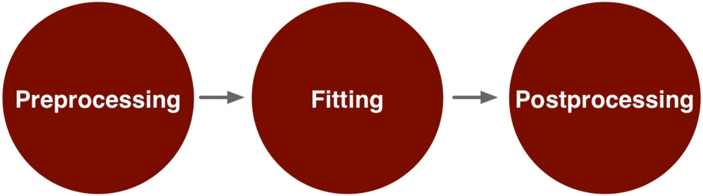

Fairness Methodology

Data transformations

Fairness Methodology

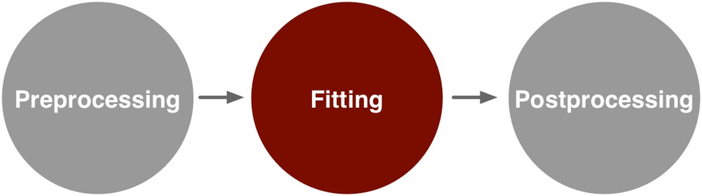

Adding variables, separate formulas, statistical learning

Fairness Methodology

Differing thresholds

Algorithmic Fairness

Typical algorithmic fairness problem in computer science has

- outcome Y
- vector X that includes a protected class or sensitive attribute $A \subset X$

Goal:

Create estimator for $f(X) = Y$ while ensuring the function is fair for A

Common measures of fairness are based on the notion of **group fairness**, striving for similarity in predicted outcomes or errors for groups

Algorithmic Fairness

Typical algorithmic fairness problem in computer science has

- outcome Y
- vector X that includes a protected class or sensitive attribute $A \subset X$

Goal:

Create estimator for $f(X) = Y$ while ensuring the function is fair for A

Common measures of fairness are based on the notion of **group fairness**, striving for similarity in predicted outcomes or errors for groups

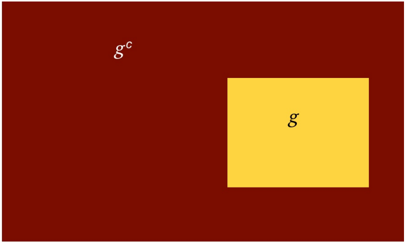

Algorithmic Fairness

Typical algorithmic fairness problem in computer science has

- outcome Y
- vector X that includes a protected class or sensitive attribute $A \subset X$

Goal:

Create estimator for $f(X) = Y$ while ensuring the function is fair for A

Common measures of fairness are based on the notion of **group fairness**, striving for similarity in predicted outcomes or errors for groups

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

- $\hat{Y}$ is predicted spending
- $\bar{Y}$ is mean spending

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

- $\hat{Y}$ is predicted spending
- $\bar{Y}$ is mean spending

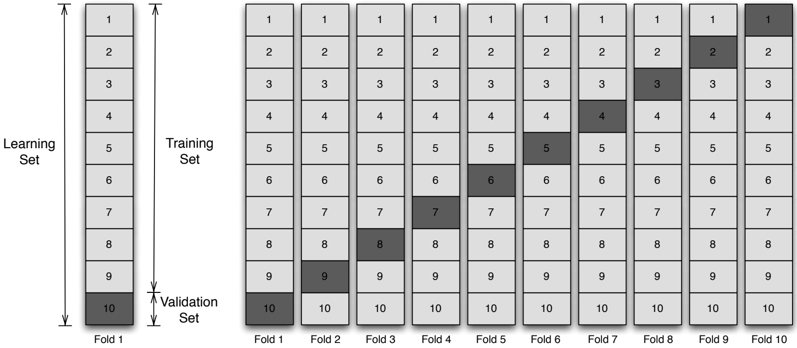

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

Health Economics

Net Compensation

(Layton et al. 2017)

$$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i)$$

Global vs. Group Fit Metrics

	R^2	MHSUD Net Compensation
1. baseline formula	13.1%	-$2,822

THE AMERICAN STATISTICIAN

Limitations of *P*-Values and *R*-Squared for Stepwise Regression Building: A Fairness Demonstration in Health Policy Risk Adjustment

Sherri Rose and Thomas G. McGuire

Global vs. Group Fit Metrics

	R^2	MHSUD Net Compensation
1. baseline formula	13.1%	-$2,822
	2% increase	31% increase
2. + mental health	13.3%	-$1,952

Limitations of *P*-Values and *R*-Squared for Stepwise Regression Building: A Fairness Demonstration in Health Policy Risk Adjustment

Sherri Rose and Thomas G. McGuire

Global vs. Group Fit Metrics

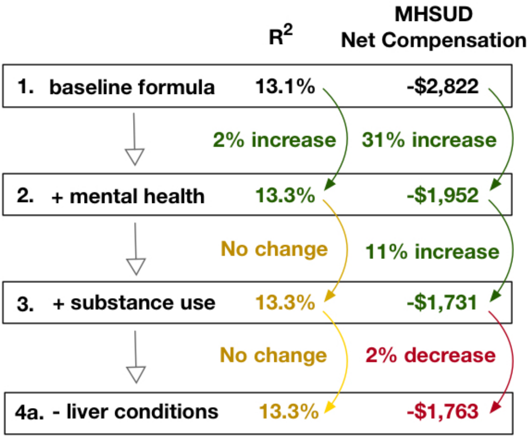

Limitations of *P*-Values and *R*-Squared for Stepwise Regression Building: A Fairness Demonstration in Health Policy Risk Adjustment

Sherri Rose and Thomas G. McGuire

Global vs. Group Fit Metrics

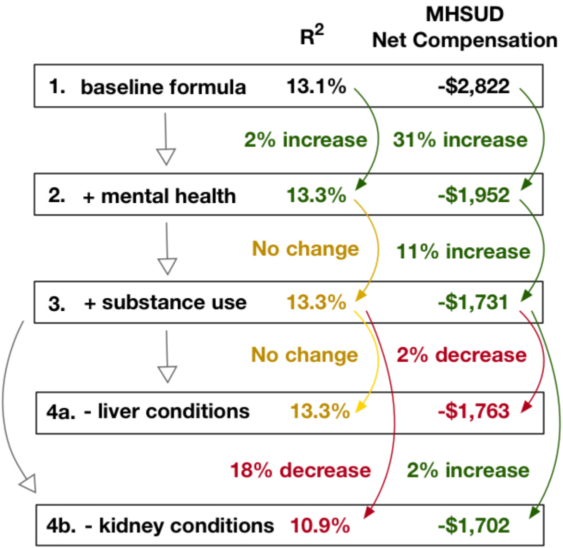

Limitations of *P*-Values and *R*-Squared for Stepwise Regression Building: A Fairness Demonstration in Health Policy Risk Adjustment

Sherri Rose and Thomas G. McGuire

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

Health Economics	Computer Science & Statistics
Net Compensation	Mean Residual Difference
(Layton et al. 2017)	(Calders et al. 2013)
$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i)$	$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i) - \frac{1}{n_c} \sum_{j \in g^c} (\hat{Y}_j - Y_j)$

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

Health Economics	Computer Science & Statistics
Net Compensation	Mean Residual Difference
(Layton et al. 2017)	(Calders et al. 2013)
$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i)$	$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i) - \frac{1}{n_c} \sum_{j \in g^c} (\hat{Y}_j - Y_j)$

Global vs. Group Fit Metrics

$$R^2 = 1 - \frac{\sum_k (Y_k - \hat{Y}_k)^2}{\sum_k (Y_k - \bar{Y}_k)^2}$$

Health Economics	Computer Science & Statistics

Health Economics

Net Compensation

(Layton et al. 2017)

$$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i)$$

Predictive Ratios

(Pope et al. 2004)

$$\frac{\sum_{i \in g} \hat{Y}_i}{\sum_{i \in g} Y_i}$$

Computer Science & Statistics

Mean Residual Difference

(Calders et al. 2013)

$$\frac{1}{n_g} \sum_{i \in g} (\hat{Y}_i - Y_i) - \frac{1}{n_c} \sum_{j \in g^c} (\hat{Y}_j - Y_j)$$

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Challenges:

- Current formulas created with parametric regression without built-in fairness criteria
- Much of the fairness literature considers binary decision-making

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Challenges:

- Current formulas created with parametric regression without built-in fairness criteria
- Much of the fairness literature considers binary decision-making

Advances:

1. Zink & Rose (2020), *Biometrics*
   Fair regression for a single attribute with continuous outcomes

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Challenges:

- Current formulas created with parametric regression without built-in fairness criteria
- Much of the fairness literature considers binary decision-making

Advances:

1. Zink & Rose (2020), *Biometrics*
   Fair regression for a single attribute with continuous outcomes
2. McGuire, Zink & Rose (2021), *American Journal of Health Economics*
   Fair regression for several single attributes and postprocessing

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Challenges:

- Current formulas created with parametric regression without built-in fairness criteria
- Much of the fairness literature considers binary decision-making

Advances:

1. Zink & Rose (2020), *Biometrics*
   Fair regression for a single attribute with continuous outcomes
2. McGuire, Zink & Rose (2021), *American Journal of Health Economics*
   Fair regression for several single attributes and postprocessing
3. Zink & Rose (2021), *BMJ Health & Care Informatics*
   Identifying complex groups defined by multiple attributes

Can we improve fairness for undercompensated groups in plan payment risk adjustment?

Challenges:

- Current formulas created with parametric regression without built-in fairness criteria
- Much of the fairness literature considers binary decision-making

Advances:

1. Zink & Rose (2020), *Biometrics*
   Fair regression for a single attribute with continuous outcomes
2. McGuire, Zink & Rose (2021), *American Journal of Health Economics*
   Fair regression for several single attributes and postprocessing
3. Zink & Rose (2021), *BMJ Health & Care Informatics*
   Identifying complex groups defined by multiple attributes

Anna Zink
PhD Student
Harvard

1 Covariance Regression

Covariance techniques require covariance between the residual and protected class be close to zero (Zafar et al. 2017a,b)

We extend these methods for **continuous residuals** with continuous Y. The new optimization problem is given by:

$$\underset{\theta}{\text{minimize}} \left\{ \sum_k \left(Y_k - \sum_p \theta_p X_{kp} \right)^2 \right\}, \quad \text{subject to}$$

$$(1 - P(A = 1)) \sum_{i \in g} \left(Y_i - \sum_p \theta_p X_{ip} \right) - P(A = 1) \sum_{j \in g^c} \left(Y_j - \sum_p \theta_p X_{jp} \right) < c,$$

where $c = m \times c^*$ with $m \in [0, 1]$ and c^* the covariance of the undercompensated group and OLS residual

1 Net Compensation Regression

Propose new custom penalty term that punishes large net compensation

Our minimization problem:

$$\sum_k \left(Y_k - \sum_p \theta_p X_{kp} \right)^2 + \lambda \left(\frac{1}{n_g} \sum_{i \in g} \left(Y_i - \sum_p \theta_p X_{ip} \right) \right)$$

① Net Compensation Regression

Propose new custom penalty term that punishes large net compensation

Our minimization problem:

$$\sum_k \left(Y_k - \sum_p \theta_p X_{kp} \right)^2 + \lambda \left(\frac{1}{n_g} \sum_{i \in g} \left(Y_i - \sum_p \theta_p X_{ip} \right) \right)$$

Can alternatively present our new method as a constraint:

$$\underset{\theta}{\text{minimize}} \left\{ \sum_k \left(Y_k - \sum_p \theta_p X_{kp} \right)^2 \right\}, \text{ subject to}$$

$$\frac{1}{n_g} \sum_{i \in g} \left(Y_i - \sum_p \theta_p X_{ip} \right) \leq z,$$

where z is positive, 1-to-1 correspondence with λ when constraint is binding

1 Large Gains in Group Fairness vs. OLS

Regression Method	R^2	MHSUD Net Compensation
Average	12.4%	
Covariance	12.4	
Net Compensation	12.5	
Weighted Average	12.6	
Mean Residual Difference	12.8	
Ordinary Least Squares	12.9	

IBS Biometrics JOURNAL OF THE INTERNATIONAL BIOMETRIC SOCIETY

Fair regression for health care spending

Anna Zink, Sherri Rose

① Large Gains in Group Fairness vs. OLS

Regression Method	R^2	MHSUD Net Compensation
Average	12.4%	
Covariance	12.4	
Net Compensation	12.5	
Weighted Average	12.6	
Mean Residual Difference	12.8	
Ordinary Least Squares	12.9	

4%

Biometrics JOURNAL OF THE INTERNATIONAL BIOMETRIC SOCIETY

Fair regression for health care spending

Anna Zink, Sherri Rose

1 Large Gains in Group Fairness vs. OLS

Regression Method	R^2	MHSUD Net Compensation
Average	12.4%	-$46
Covariance	12.4	-46
Net Compensation	12.5	-232
Weighted Average	12.6	-411
Mean Residual Difference	12.8	-1208
Ordinary Least Squares	12.9	-1872

98%

Biometrics JOURNAL OF THE INTERNATIONAL BIOMETRIC SOCIETY

Fair regression for health care spending

Anna Zink, Sherri Rose

2 Multiple Groups

Improving the Performance of Risk Adjustment Systems: Constrained Regressions, Reinsurance, and Variable Selection

Thomas G. McGuire, Anna L. Zink and Sherri Rose

AMERICAN JOURNAL OF HEALTH ECONOMICS

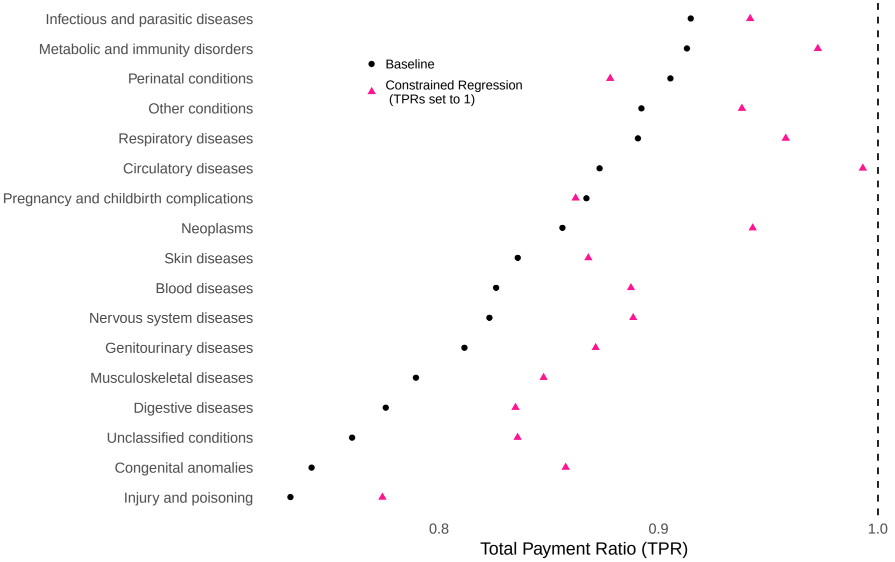

③ Complex Groups

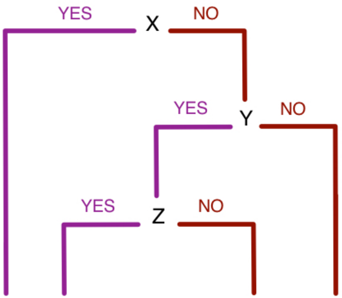

③ Complex Groups

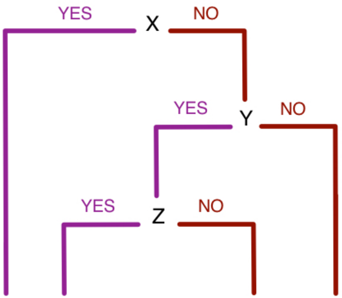

Example Hypothetical Group

3 Complex Groups

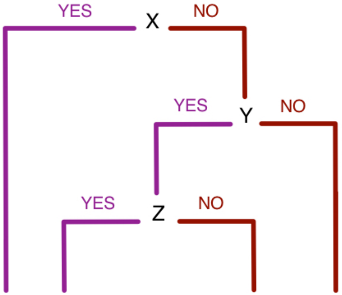

Example Hypothetical Group

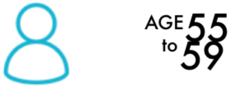

3 Complex Groups

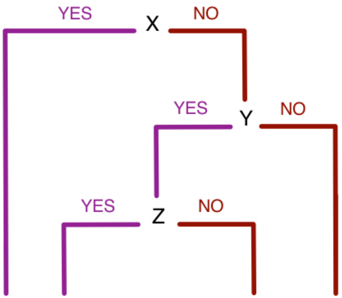

Example Hypothetical Group

3 Complex Groups

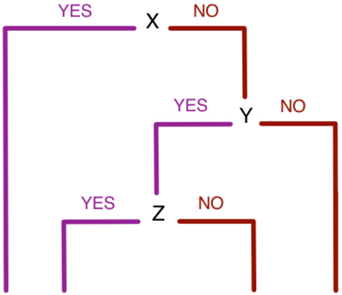

Example Hypothetical Group

Undercompensated?

3 Complex Groups

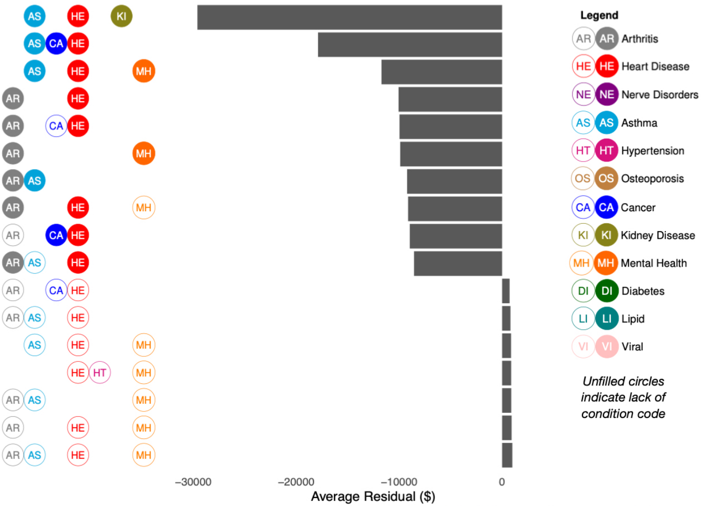

Identifying undercompensated groups defined by multiple attributes in risk adjustment

Anna Zink, Sherri Rose

BMJ Health & Care Informatics

IN CLOSING

Biases enter data & algorithms in many ways

Biases enter data & algorithms in many ways

Diverse teams

Biases enter data & algorithms in many ways

Diverse teams

Metric<u>s</u> matter

Biases enter data & algorithms in many ways

Diverse teams

Metrics matter

Not as simple as add or drop attribute

Biases enter data & algorithms in many ways

Diverse teams

Metrics matter

Not as simple as add or drop attribute

Algorithms may contribute to solutions

Biases enter data & algorithms in many ways

Diverse teams

Metrics matter

Not as simple as add or drop attribute

Algorithms may contribute to solutions

Respect the data

Biases enter data & algorithms in many ways

Diverse teams

Metrics matter

Not as simple as add or drop attribute

Algorithms may contribute to solutions

Respect the data

Engage with the application or do not use it

Biases enter data & algorithms in many ways

Diverse teams

Metrics matter

Not as simple as add or drop attribute

Algorithms may contribute to solutions

Respect the data

Engage with the application or do not use it

Cite the literature

Does Your Algorithm Have a Social Impact Statement?

Responsibility

Explainability

Accuracy

Auditability

Fairness

fatml.org/resources/principles-for-accountable-algorithms

If you don't meet people like you in your courses or see yourself in your instructors, that doesn't mean you don't belong in this field

Acknowledgements


Sam Adhikari, PhD
NYU


Austin Denteh, PhD
Tulane


Savannah Bergquist, PhD
Berkeley Haas


Akritee Shrestha, MS
Wayfair


Maia Majumder, PhD
Boston Children's/Harvard


Alex McDowell, PhD
MGH/Harvard


Toyya Pujol, PhD
Purdue


Irina Degtiar, PhD
Mathematica


Anna Zink
Harvard


Paidamoyo Chapfuwa, PhD
Stanford


Samson Mataraso
Stanford


Checo Gonzales
Stanford

Funding:

NIH Director's New Innovator Award (DP2-MD012722)
Laura and John Arnold Foundation