

Teaching "Responsible Data Science"

Debbie Yuster, Ph.D.


Assistant Professor of Data Science and Mathematics
Ramapo College of New Jersey
dyuster@ramapo.edu

Slides: bit.ly/yuster-mathfest-22

About me



(Project NEXt Sun Dot 2007)



**RAMAPO
COLLEGE**
OF NEW JERSEY

THE WALL STREET JOURNAL.

Slides: bit.ly/yuster-mathfest-22

DATA 620: Ethics in Data and Computing

Required for students in:

- MS Data Science (2020)
- MS Applied Mathematics (2022)
- MS Computer Science (2022)


Prerequisite for thesis and fieldwork



NYU's Responsible Data Science course

Created by **Julia Stoyanovich** (NYU Associate Professor of Data Science, Computer Science and Engineering)

With involvement from **George Wood** (NYU Moore-Sloan Faculty Fellow, Center for Data Science)



Ethics course for *techies*

- Rigorous mathematics - yes, really!
- Python Jupyter notebooks running state of the art packages

But also-


- Social Science
- Philosophy
- Management
- Legal issues

Sample materials

The evils of discrimination

Disparate treatment
is the illegal practice of treating an entity, such as a job applicant or an employee, differently based on a **protected characteristic** such as race, gender, age, disability status, religion, sexual orientation, or national origin.

Disparate impact
is the result of systematic disparate treatment, where disproportionate **adverse impact** is observed on members of a **protected class**.

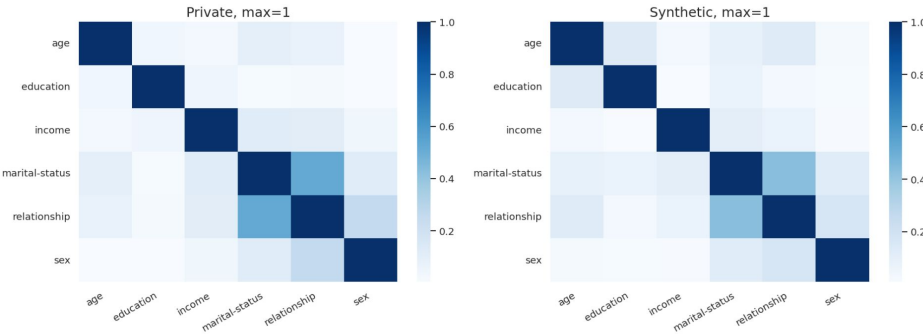


Slide

We can even plot the mutual information between all pairs of features in the dataset:

```
[ ] attribute_description = read_json_file(description_files['correlated_attribute_mode'])['attribute_description']
ModelInspector(real_data,
               synthetic_correlated,
               attribute_description).mutual_information_heatmap()
```

Pairwise Mutual Information Comparison (Private vs Synthetic)



Feature 1 \ Feature 2	Private, max=1	Synthetic, max=1
age	0.95	0.85
education	0.85	0.75
income	0.75	0.65
marital-status	0.65	0.55
relationship	0.55	0.45
sex	0.45	0.35

Lab

Course outline

1. Fairness

- Sociotechnical systems; Stakeholders; Trade-offs
- Fairness in classification; Fairness in risk assessment
- Causality; Counterfactual fairness
- Philosophy of fairness

2. Data Science Lifecycle

- Data profiling and cleaning
- Taming technical bias

3. Data Protection

- Limits of anonymization
- Differential privacy; Synthetic data
- Applied Ethics; Ethical frameworks

4. Transparency and Interpretability

- Auditing black-box models
- Explainable machine learning
- Online ad delivery
- Legal frameworks

Materials available (by request)

- Week by week reading list
- Lecture slides
- Lab activities
- Homework assignments and solutions
- Final project or paper instructions

Resources

- Spring 2022 Responsible Data Science course website:
<https://dataresponsibly.github.io/rds/>
- Journal article: <https://rdcu.be/cS0h6>

Lewis, A. and Stoyanovich, J. 2021. Teaching Responsible Data Science: Charting New Pedagogical Territory. *International Journal of Artificial Intelligence in Education*. (2021), 1–25. DOI:<https://doi.org/10.1007/s40593-021-00241-7>.

- Light-ish background reading:
The Ethical Algorithm by Michael Kearns & Aaron Roth
- dyuster@ramapo.edu and/or stoyanovich@nyu.edu

