DS-GA 3001.009: Responsible Data Science

Interpretability

Prof. Julia Stoyanovich
Center for Data Science
Computer Science and Engineering at Tandon

@stoyanoj

http://stoyanovich.org/
https://dataresponsibly.github.io/

What is interpretability?

- **Explaining black-box models**
- **Online ad targeting**
- **Interpretability**

A kitchen sink? Or a foundational concept for responsible data science?

https://favpng.com/png_view/cartoon-kitchen-sink-scene-towel-sink-kitchen-cartoon-png/pMFrA1n9

Algorithmic rankers

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

Input: database of items (individuals, colleges, cars, …)

Score-based ranker: computes the score of each item using a **known formula**, often a monotone aggregation function, then sorts items on score

Output: permutation of the items, complete or top-k

$\mathcal{D}$			f
id	x_1	x_2	$x_1 + x_2$
t_1	0.63	0.71	1.34
t_2	0.72	0.65	1.37
t_3	0.58	0.78	1.36
t_4	0.7	0.68	1.38
t_5	0.53	0.82	1.35
t_6	0.61	0.79	1.4

Do we have transparency?

We have syntactic transparency, but lack interpretability!

Opacity in algorithmic rankers

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

Reason 1: The scoring formula alone does not indicate the relative rank of an item.

Scores are absolute, rankings are relative. Is 5 a good score? What about 10? 15?

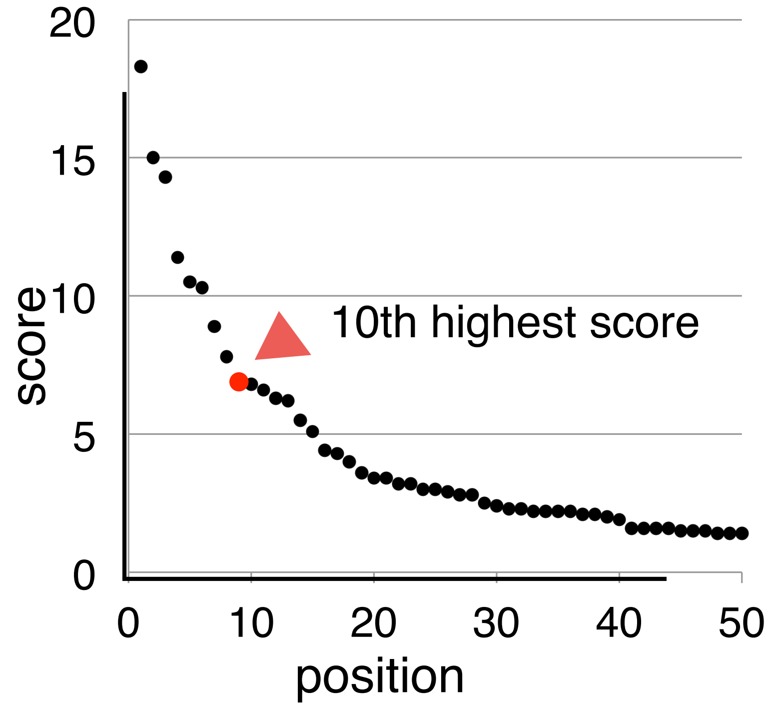

Opacity in algorithmic rankers

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

Reason 2: A ranking may be unstable if there are tied or nearly-tied items.

Rank	Institution	Average Count	Faculty
1	▶ Carnegie Mellon University	18.4	123
2	▶ Massachusetts Institute of Technology	15.6	64
3	▶ Stanford University	14.8	56
4	▶ University of California - Berkeley	11.5	50
5	▶ University of Illinois at Urbana-Champaign	10.6	56
6	▶ University of Washington	10.3	50
7	▶ Georgia Institute of Technology	8.9	81
8	▶ University of California - San Diego	8	51
9	▶ Cornell University	7	45
10	▶ University of Michigan	6.8	63
11	▶ University of Texas - Austin	6.6	43
12	▶ University of Massachusetts - Amherst	6.4	47

Opacity in algorithmic rankers

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

Reason 3: A ranking methodology may be unstable: small changes in weights can trigger significant re-shuffling.

THE NEW YORKER

DEPT. OF EDUCATION FEBRUARY 14 & 21, 2011 ISSUE

THE ORDER OF THINGS

What college rankings really tell us.

By Malcolm Gladwell

1. Porsche Cayman 193
2. Chevrolet Corvette 186
3. Lotus Evora 182

1. Chevrolet Corvette 205
2. Lotus Evora 195
3. Porsche Cayman 195

1. Lotus Evora 205
2. Porsche Cayman 198
3. Chevrolet Corvette 192

https://www.newyorker.com/magazine/2011/02/14/the-order-of-things

Opacity in algorithmic rankers

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

Reason 4: The **weight of an attribute** in the scoring formula **does not determine its impact** on the outcome.

Rank	Name	Avg Count	Faculty	Pubs	GRE
1	CMU	18.3	122	2	791
2	MIT	15	64	3	772
3	Stanford	14.3	55	5	800
4	UC Berkeley	11.4	50	3	789
5	UIUC	10.5	55	3	772
6	UW	10.3	50	2	796
		….			
39	U Chicago	2	28	2	779
40	UC Irvine	1.9	28	2	787
41	BU	1.6	15	2	783
41	U Colorado Boulder	1.6	32	1	761
41	UNC Chapel Hill	1.6	22	2	794
41	Dartmouth	1.6	18	2	794

Given a score function:

$$0.2 * faculty +$$
$$0.3 * avg\ cnt +$$
$$0.5 * gre$$

Rankings are not benign!

THE NEW YORKER

DEPT. OF EDUCATION FEBRUARY 14 & 21, 2011 ISSUE

THE ORDER OF THINGS

What college rankings really tell us.

By Malcolm Gladwell

Rankings are not benign. They enshrine very particular ideologies, and, at a time when American higher education is facing a crisis of accessibility and affordability, we have adopted **a de-facto standard of college quality** that is uninterested in both of those factors. And why? Because a group of magazine analysts in an office building in Washington, D.C., decided twenty years ago to **value selectivity over efficacy**, to **use proxies** that scarcely relate to what they're meant to be proxies for, and to **pretend that they can compare** a large, diverse, low-cost land-grant university in rural Pennsylvania with a small, expensive, private Jewish university on two campuses in Manhattan.

Interpretability in the service of trust!

Gladwell makes the point that rankings are claiming objectivity, yet are comparing apples and oranges.

In that sense, **a score-based ranker is a quintessential "black box" of data science**, and perhaps the simplest possible such black box.

AI is a red herring, privacy / IP / gaming arguments are overused. The truly difficult issues are that:

1) using math to pretend that we are correct when making intrinsically subjective decisions reinforcing the balance of power in society
2) that math / objectivity is used as a substitute for trust, but **trust must run deeper than math**!
3) need to find a kind of an interpretability that will enable trust!

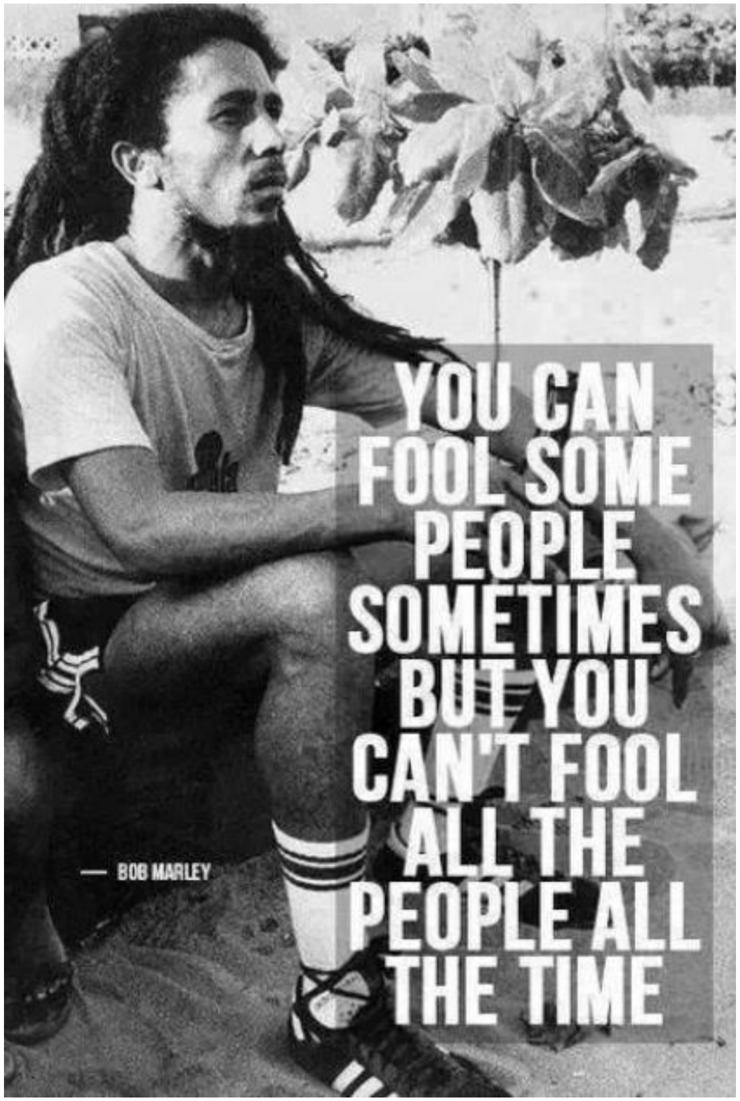

Harms of opacity

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

1. **Due process / fairness.** The subjects of the ranking cannot have confidence that their ranking is meaningful or correct, or that they have been treated like similarly situated subjects - **procedural regularity**

2. **Hidden normative commitments.** What factors does the vendor encode in the scoring ranking process? What are the **actual** effects of the scoring / ranking process? Is it stable? How was it validated?

Harms of opacity

https://freedom-to-tinker.com/2016/08/05/revealing-algorithmic-rankers/

3. **Interpretability.** Especially where ranking algorithms are performing a public function, **political legitimacy** requires that the public be able to interpret algorithmic outcomes in a meaningful way. Avoid *algocracy*: the rule by incontestable algorithms.

4. **Meta-methodological assessment.** Is *a* ranking / *this* ranking appropriate here? Can we use a process if it cannot be explained? Probably yes, for recommending movies. Probably not for college admissions.

an (ongoing) attempt at regulation

New York City Local Law 49

January 11, 2018

Local Law 49 of 2018 in relation to automated decision systems used by agencies

THE NEW YORK CITY COUNCIL
Corey Johnson, Speaker

Sign In

LEGISLATIVE RESEARCH CENTER

Council Home | Legislation | Calendar | City Council | Committees

RSS | Alerts

Details | Reports

File #:	Int 1696-2017 Version: A	Name:	Automated decision systems used by agencies.
Type:	Introduction	Status:	Enacted
		Committee:	Committee on Technology
On agenda:	8/24/2017		
Enactment date:	1/11/2018	Law number:	2018/049
Title:	A Local Law in relation to automated decision systems used by agencies		
Sponsors:	James Vacca, Helen K. Rosenthal, Corey D. Johnson, Rafael Salamanca, Jr., Vincent J. Gentile, Robert E. Cornegy, Jr., Jumaane D. Williams, Ben Kallos, Carlos Menchaca		
Council Member Sponsors:	9		
Summary:	This bill would require the creation of a task force that provides recommendations on how information on agency automated decision systems may be shared with the public and how agencies may address instances where people are harmed by agency automated decision systems.		
Indexes:	Oversight		
Attachments:	1. Summary of Int. No. 1696-A, 2. Summary of Int. No. 1696, 3. Int. No. 1696, 4. August 24, 2017 - Stated Meeting Agenda with Links to Files, 5. Committee Report 10/16/17, 6. Hearing Testimony 10/16/17, 7. Hearing Transcript 10/16/17, 8. Proposed Int. No. 1696-A - 12/12/17, 9. Committee Report 12/7/17, 10. Hearing Transcript 12/7/17, 11. December 11, 2017 - Stated Meeting Agenda with Links to Files, 12. Hearing Transcript - Stated Meeting 12-11-17, 13. Int. No. 1696-A (FINAL), 14. Fiscal Impact Statement, 15. Legislative Documents - Letter to the Mayor, 16. Local Law 49, 17. Minutes of the Stated Meeting - December 11, 2017		

The original draft

August 16, 2017

Int. No. 1696

By Council Member Vacca

A Local Law to amend the administrative code of the city of New York, in relation to automated processing of data for the purposes of targeting services, penalties, or policing to persons

Be it enacted by the Council as follows:

Section 1. Section 23-502 of the administrative code of the city of New York is amended to add a new subdivision g to read as follows:

g. Each agency that uses, for the purposes of targeting services to persons, imposing penalties upon persons or policing, an algorithm or any other method of automated processing system of data shall:

1. Publish on such agency's website, the source code of such system; and

2. Permit a user to (i) submit data into such system for self-testing and (ii) receive the results of having such data processed by such system.

§ 2. This local law takes effect 120 days after it becomes law.

MAJ
LS# 10948
8/16/17 2:13 PM

not what was adopted

How I got involved

October 16, 2017

THE NEW YORKER

By Julia Powles December 20, 2017

ELEMENTS

NEW YORK CITY'S BOLD, FLAWED ATTEMPT TO MAKE ALGORITHMS ACCOUNTABLE

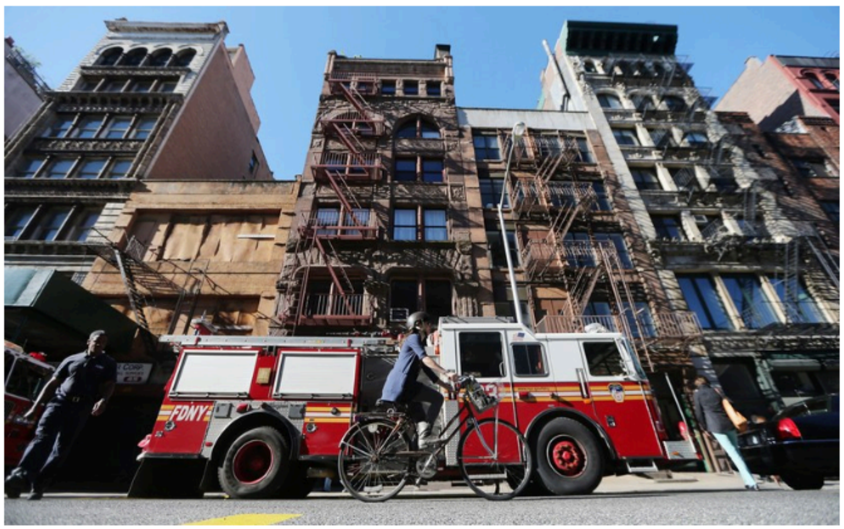

Automated systems guide the allocation of everything from firehouses to food stamps. So why don't we know more about them?

Photograph by Mario Tama / Getty

https://dataresponsibly.github.io/documents/Stoyanovich_VaccaBill.pdf

Summary of Local Law 49

January 11, 2018

An **Automated Decision System (ADS)** is a "computerized implementation of algorithms, including those derived from machine learning or other data processing or artificial intelligence techniques, which are used to make or assist in making decisions."

Form task force that surveys the current use of ADS in City agencies and develops procedures for:

- requesting and receiving an **explanation** of an algorithmic decision affecting an individual (3(b))
- interrogating ADS for **bias and discrimination** against members of legally-protected groups (3(c) and 3(d))
- allowing the **public** to **assess** how ADS function and are used (3(e)), and archiving ADS together with the data they use (3(f))

The ADS Task Force

May 16, 2018

Visit **alpha.nyc.gov** to help us test out new ideas for NYC's website.

The Official Website of the City of New York

NYC

简体中文 ▸ Translate | ▼ Text Size

NYC Resources | NYC311 | Office of the Mayor | Events | Connect | Jobs | Search

Mayor | First Lady | News | Officials

SHARE

Email

Print

Mayor de Blasio Announces First-In-Nation Task Force To Examine Automated Decision Systems Used By The City

May 16, 2018

NEW YORK— Today, Mayor de Blasio announced the creation of the Automated Decision Systems Task Force which will explore how New York City uses algorithms. The task force, the first of its kind in the U.S., will work to develop a process for reviewing "automated decision systems," commonly known as algorithms, through the lens of equity, fairness and accountability.

"As data and technology become more central to the work of city government, the algorithms we use to aid decision making must be aligned with our goals and values," said **Mayor de Blasio**. "The establishment of the Automated Decision Systems Task Force is an important first step towards greater transparency and equity in our use of technology."

The outcome (so far)

November 19, 2019

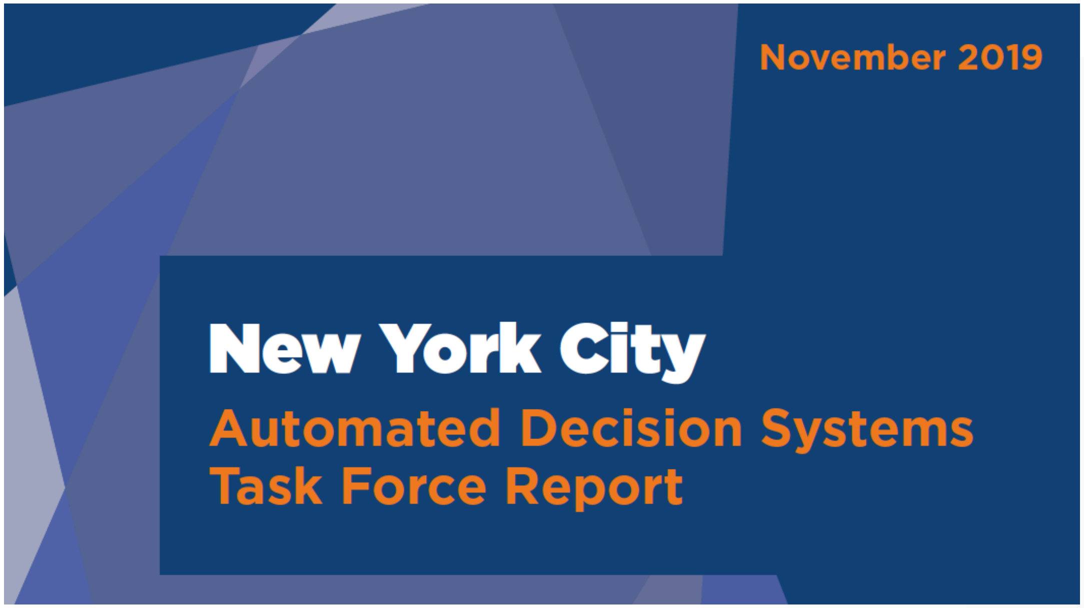

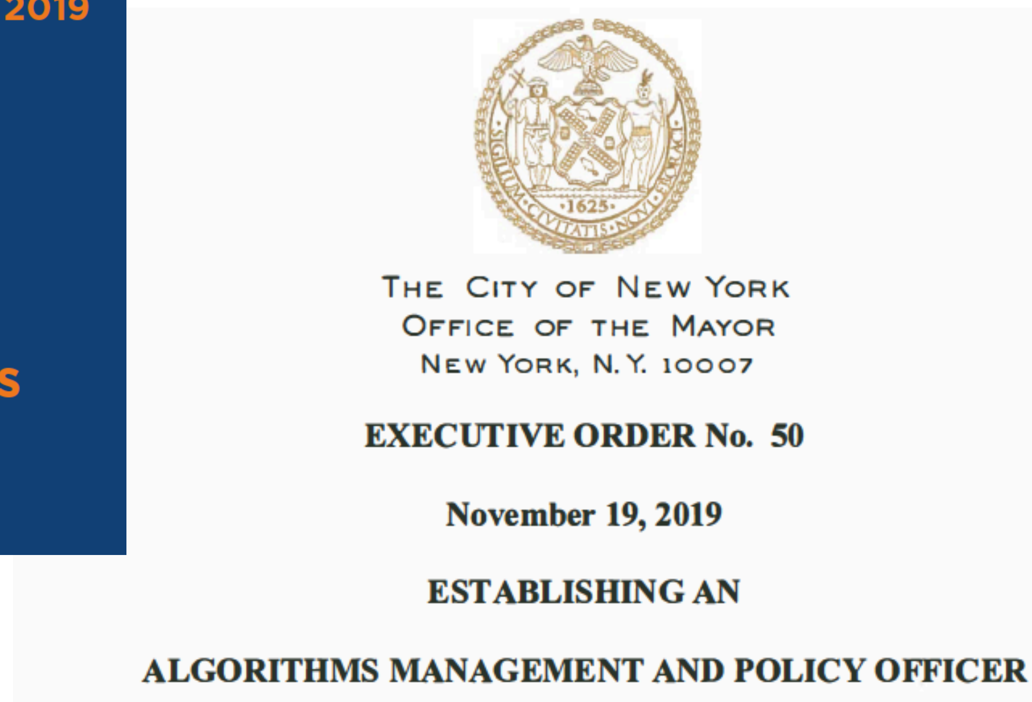

THE CITY OF NEW YORK
OFFICE OF THE MAYOR
NEW YORK, N.Y. 10007

EXECUTIVE ORDER No. 50

November 19, 2019

ESTABLISHING AN

ALGORITHMS MANAGEMENT AND POLICY OFFICER

https://www1.nyc.gov/site/adstaskforce/index.page

https://www1.nyc.gov/assets/adstaskforce/downloads/pdf/ADS-Report-11192019.pdf

https://www1.nyc.gov/assets/home/downloads/pdf/executive-orders/2019/eo-50.pdf

from transparency to interpretability

Point 1

algorithmic transparency is not synonymous with releasing the source code

publishing source code helps, but it is sometimes unnecessary and often insufficient

Point 2

algorithmic transparency requires data transparency

data is used in training, validation, deployment

validity, accuracy, applicability can only be understood in the data context

data transparency is necessary for all ADS, not only for ML-based systems

Point 3

data transparency is not synonymous with making all data public

release data whenever possible;

also release:

data selection, collection and pre-processing methodologies; data provenance and quality information; known sources of bias; privacy-preserving statistical summaries of the data

Data Synthesizer

[Ping, Stoyanovich, Howe **SSDBM 2017**]

http://demo.dataresponsibly.com/synthesizer/

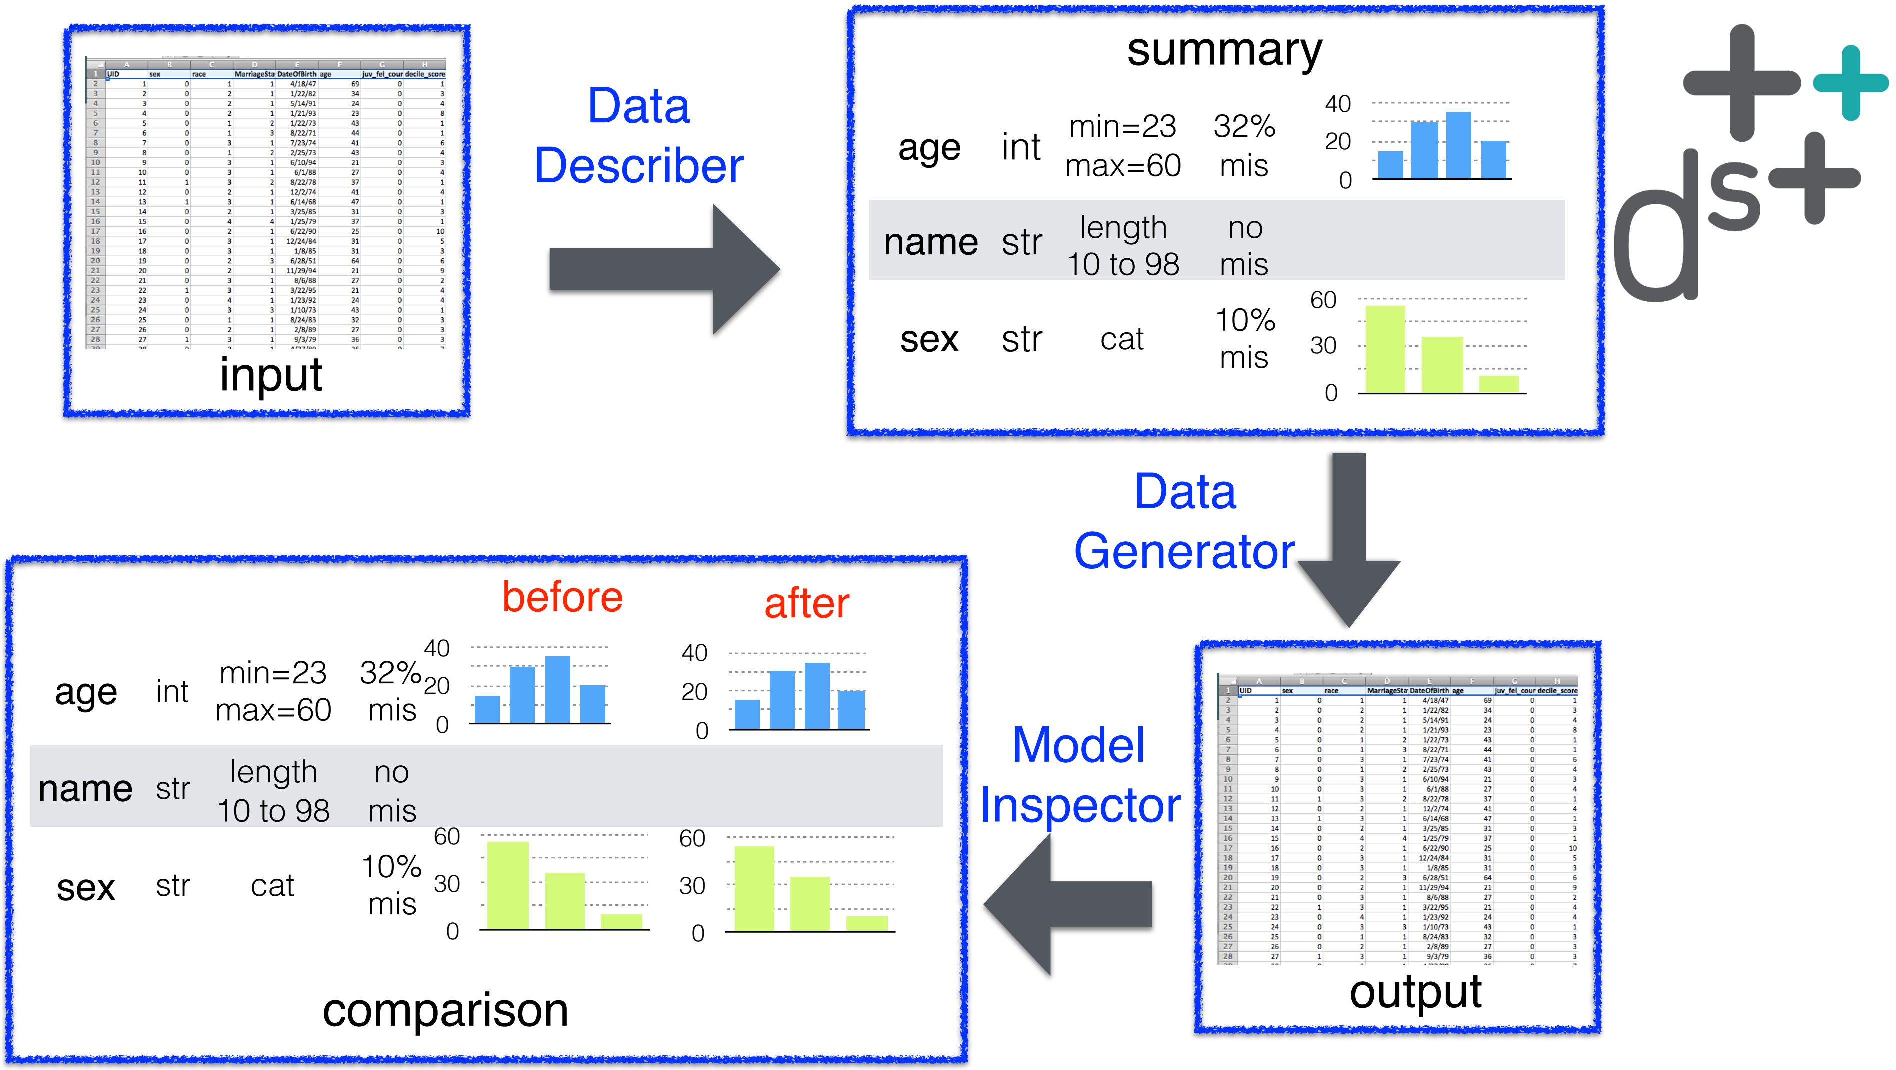

Point 4

actionable transparency requires
interpretability

explain assumptions and effects, not details of operation

engage the public - technical and non-technical

“Nutritional labels” for data and models

[K. Yang, J. Stoyanovich, A. Asudeh, B. Howe, HV Jagadish, G. Miklau; *SIGMOD 2018*]

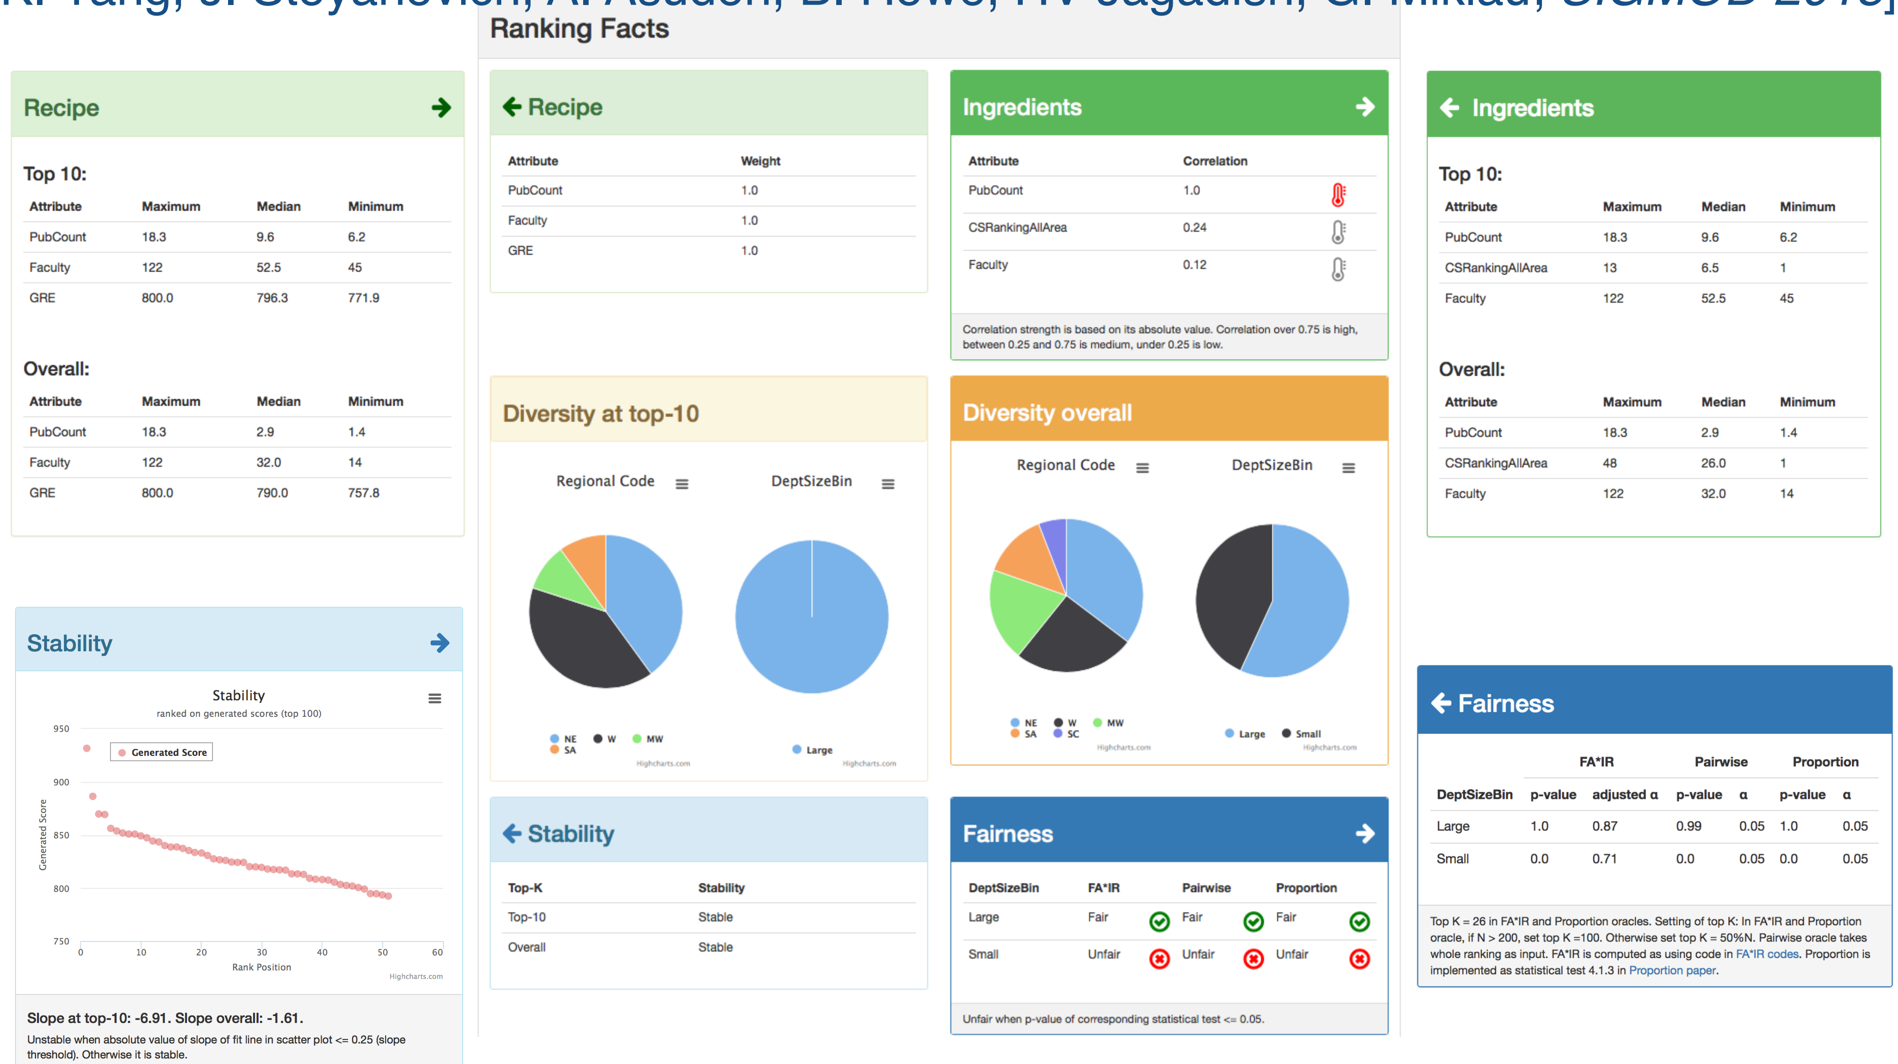

http://demo.dataresponsibly.com/rankingfacts/nutrition_facts/

Properties of a nutritional label

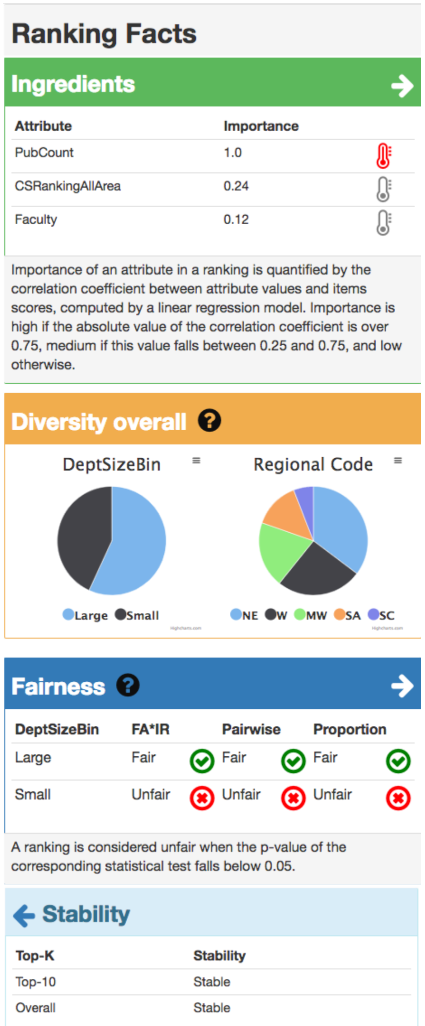

comprehensible: short, simple, clear

consultative: provide actionable info

comparable: implying a standard

concrete: helps determine a dataset's fitness for use for a given task

joint with Howe [UW] - [Data Engineering Bulletin, 2019]

Point 5

transparency by design, not as an afterthought

provision for transparency and interpretability at every stage of the data lifecycle

useful internally during development, for communication and coordination between agencies, and for accountability to the public

Frog’s eye view

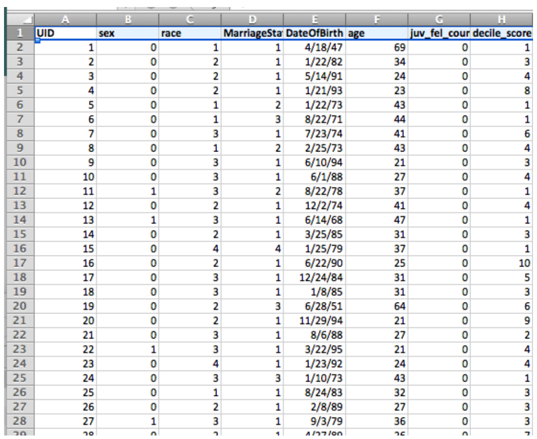

	UID	sex	race	MarriageSta	DateOfBirth	age	juv_fel_cour	decile_score
2	1	0	1	1	4/18/47	69	0	1
3	2	0	2	1	1/22/82	34	0	3
4	3	0	2	1	5/14/91	24	0	4
5	4	0	2	1	1/21/93	23	0	8
6	5	0	1	2	1/22/73	43	0	1
7	6	0	1	3	8/22/71	44	0	1
8	7	0	3	1	7/23/74	41	0	6
9	8	0	1	2	2/25/73	43	0	4
10	9	0	3	1	6/10/94	21	0	3
11	10	0	3	1	6/1/88	27	0	4
12	11	1	3	2	8/22/78	37	0	1
13	12	0	2	1	12/2/74	41	0	4
14	13	1	3	1	6/14/68	47	0	1
15	14	0	2	1	3/25/85	31	0	3
16	15	0	4	4	1/25/79	37	0	1
17	16	0	2	1	6/22/90	25	0	10
18	17	0	3	1	12/24/84	31	0	5
19	18	0	3	1	1/8/85	31	0	3
20	19	0	2	3	6/28/51	64	0	6
21	20	0	2	1	11/29/94	21	0	9
22	21	0	3	1	8/6/88	27	0	2
23	22	1	3	1	3/22/95	21	0	4
24	23	0	4	1	1/23/92	24	0	4
25	24	0	3	3	1/10/73	43	0	1
26	25	0	1	1	8/24/83	32	0	3
27	26	0	2	1	2/8/89	27	0	3
28	27	1	3	1	9/3/79	36	0	3

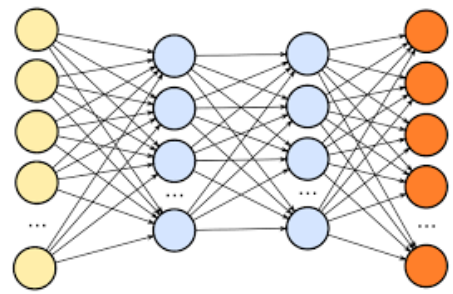

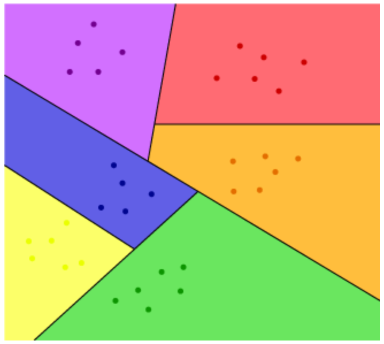

but where does the data come from?

The data science lifecycle

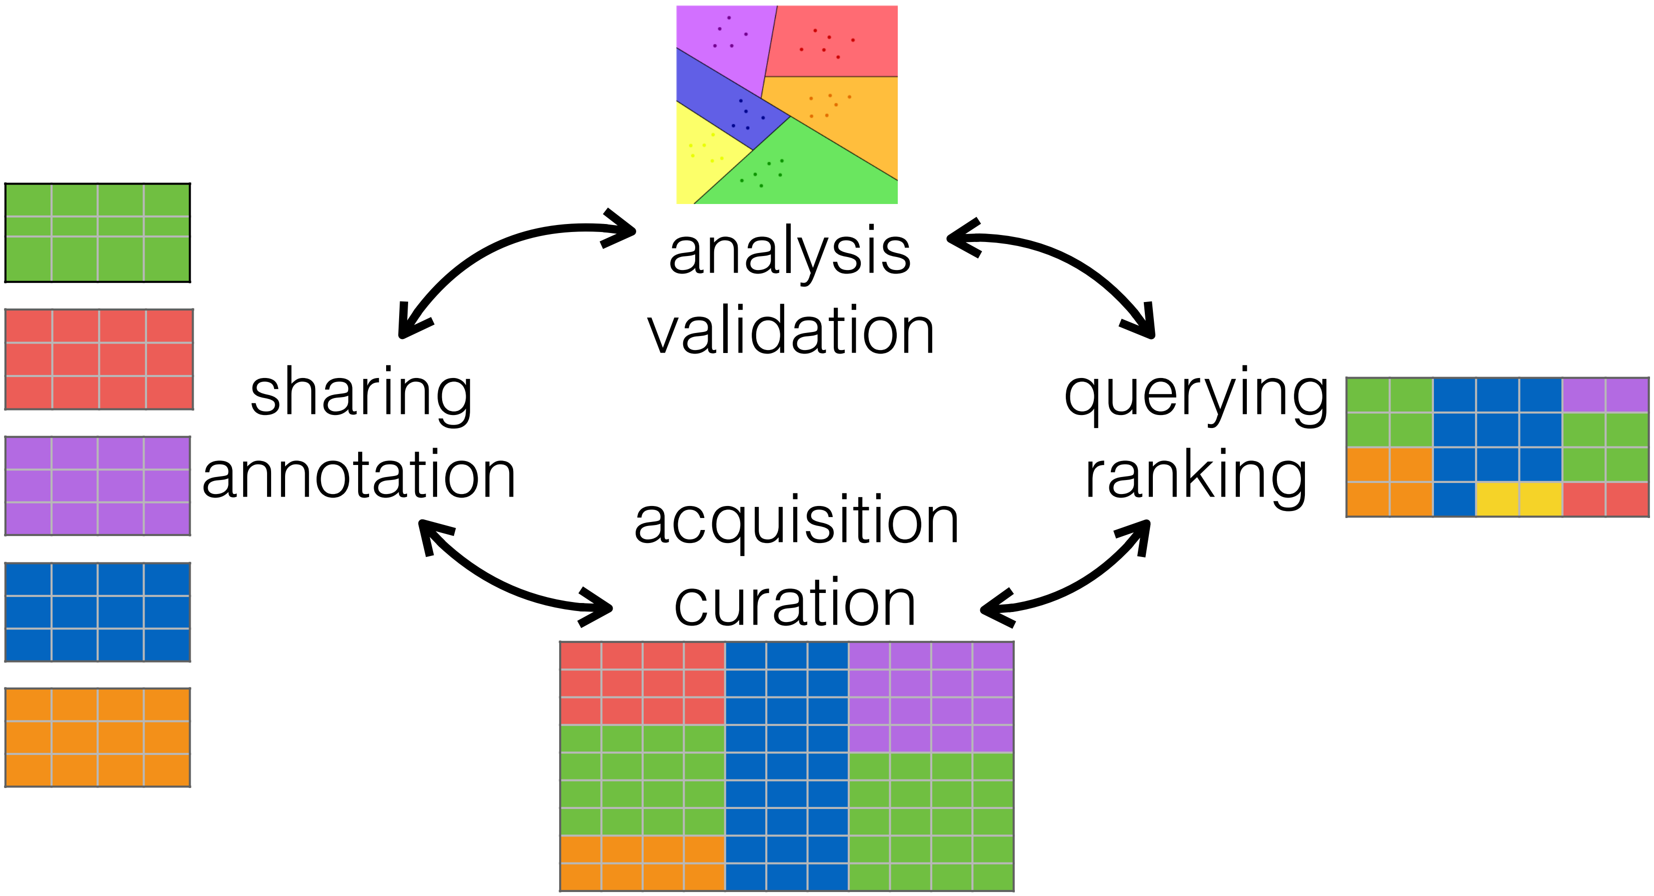

responsible data science requires a holistic view of the data lifecycle

Responsibility by design

Fides

Sharing and Curation
Annotation
Anonymization

Integration
Triage
Alignment
Transformation

Processing
Querying
Ranking
Analytics

Verification and compliance
Provenance
Explanations

Systems support for responsible data science

Responsibility by design, managed at all stages of the lifecycle of data-intensive applications

ADS example

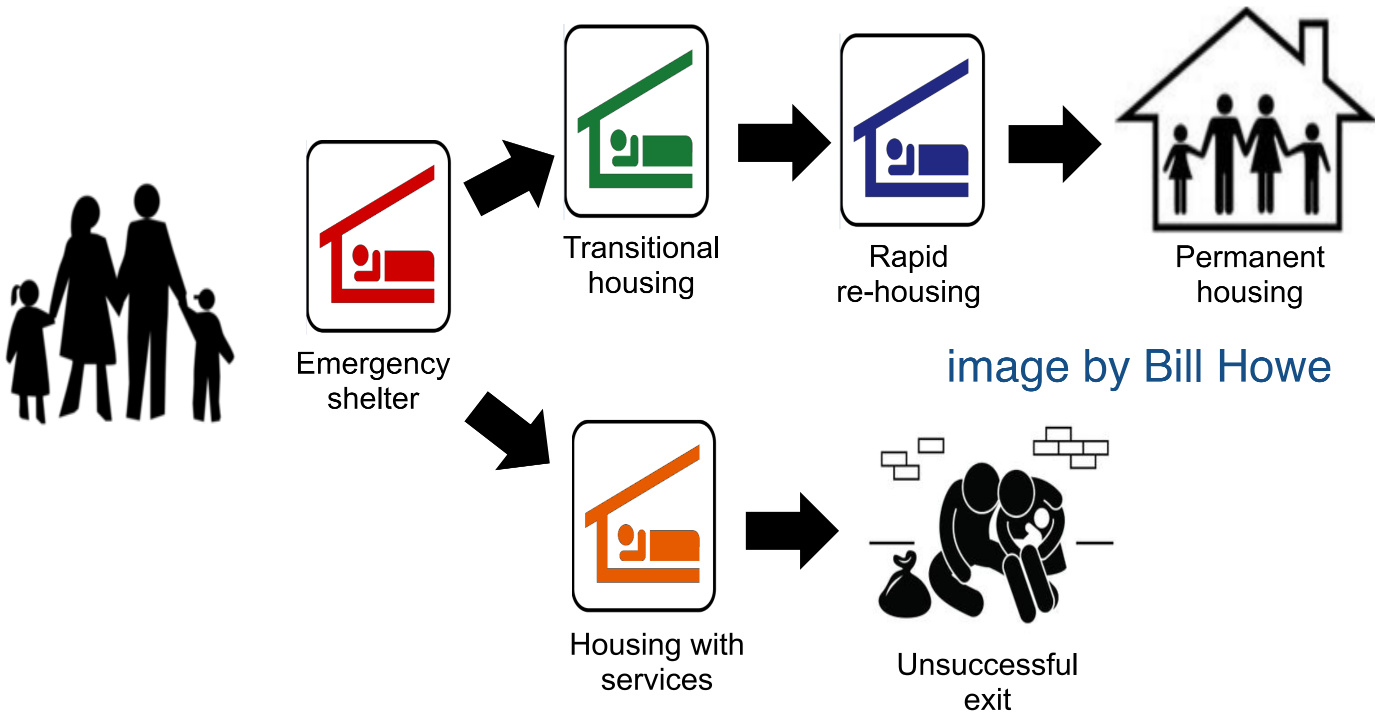

- **Allocate** interventions: services and support mechanisms
- **Recommend** pathways through the system
- **Evaluate** effectiveness of interventions, pathways, over-all system

Mitigating urban homelessness

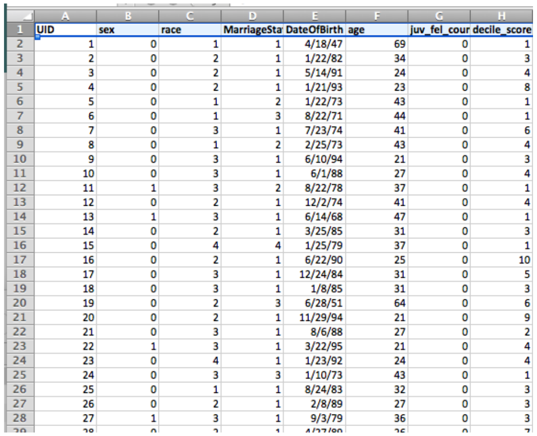

	UID	sex	race	MarriageSta	DateOfBirth	age	juv_fel_cour	decile_score
2	1	0	1	1	4/18/47	69	0	1
3	2	0	2	1	1/22/82	34	0	3
4	3	0	2	1	5/14/91	24	0	4
5	4	0	2	1	1/21/93	23	0	8
6	5	0	1	2	1/22/73	43	0	1
7	6	0	1	3	8/22/71	44	0	1
8	7	0	3	1	7/23/74	41	0	6
9	8	0	1	2	2/25/73	43	0	4
10	9	0	3	1	6/10/94	21	0	3
11	10	0	3	1	6/1/88	27	0	4
12	11	1	3	2	8/22/78	37	0	1
13	12	0	2	1	12/2/74	41	0	4
14	13	1	3	1	6/14/68	47	0	1
15	14	0	2	1	3/25/85	31	0	3
16	15	0	4	4	1/25/79	37	0	1
17	16	0	2	1	6/22/90	25	0	10
18	17	0	3	1	12/24/84	31	0	5
19	18	0	3	1	1/8/85	31	0	3
20	19	0	2	3	6/28/51	64	0	6
21	20	0	2	1	11/29/94	21	0	9
22	21	0	3	1	8/6/88	27	0	2
23	22	1	3	1	3/22/95	21	0	4
24	23	0	4	1	1/23/92	24	0	4
25	24	0	3	3	1/10/73	43	0	1
26	25	0	1	1	8/24/83	32	0	3
27	26	0	2	1	2/8/89	27	0	3
28	27	1	3	1	9/3/79	36	0	3

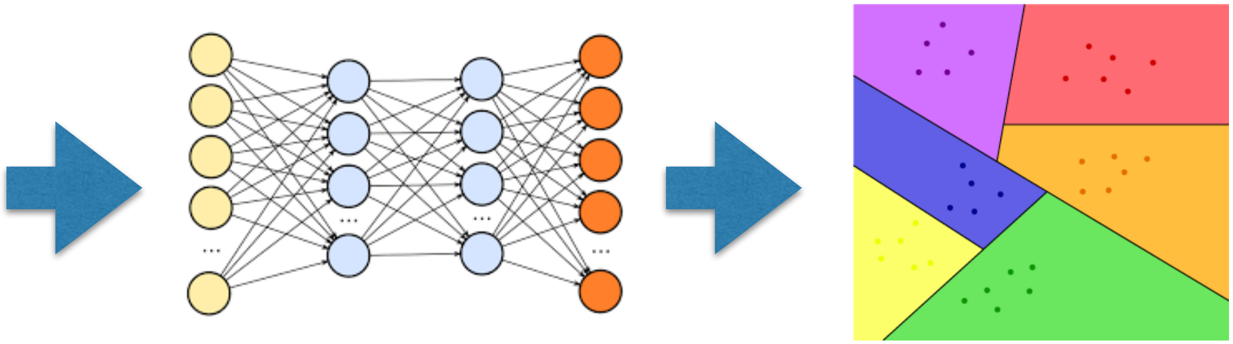

finding: women are underrepresented in the favorable outcome groups (group fairness)

fix the model!

of course, but maybe… the input was generated with:

select * from R
where status = ‘unsheltered’
and length > 2 1 month

10% female

40% female

Mitigating urban homelessness

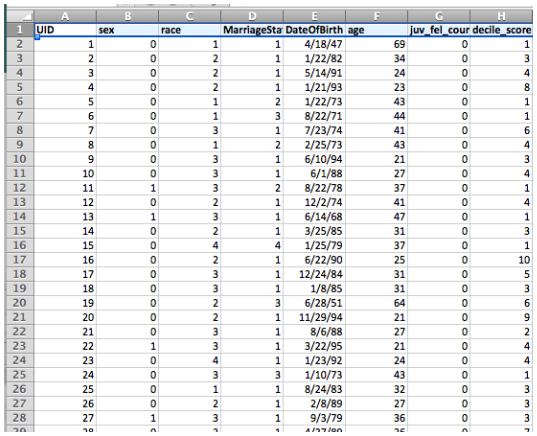

	A	B	C	D	E	F	G	H
1	UID	sex	race	MarriageSta	DateOfBirth	age	juv_fel_cour	decile_score
2	1	0	1	1	4/18/47	69	0	1
3	2	0	2	1	1/22/82	34	0	3
4	3	0	2	1	5/14/91	24	0	4
5	4	0	2	1	1/21/93	23	0	8
6	5	0	1	2	1/22/73	43	0	1
7	6	0	1	3	8/22/71	44	0	1
8	7	0	3	1	7/23/74	41	0	6
9	8	0	1	2	2/25/73	43	0	4
10	9	0	3	1	6/10/94	21	0	3
11	10	0	3	1	6/1/88	27	0	4
12	11	1	3	2	8/22/78	37	0	1
13	12	0	2	1	12/2/74	41	0	4
14	13	1	3	1	6/14/68	47	0	1
15	14	0	2	1	3/25/85	31	0	3
16	15	0	4	4	1/25/79	37	0	1
17	16	0	2	1	6/22/90	25	0	10
18	17	0	3	1	12/24/84	31	0	5
19	18	0	3	1	1/8/85	31	0	3
20	19	0	2	3	6/28/51	64	0	6
21	20	0	2	1	11/29/94	21	0	9
22	21	0	3	1	8/6/88	27	0	2
23	22	1	3	1	3/22/95	21	0	4
24	23	0	4	1	1/23/92	24	0	4
25	24	0	3	3	1/10/73	43	0	1
26	25	0	1	1	8/24/83	32	0	3
27	26	0	2	1	2/8/89	27	0	3
28	27	1	3	1	9/3/79	36	0	3

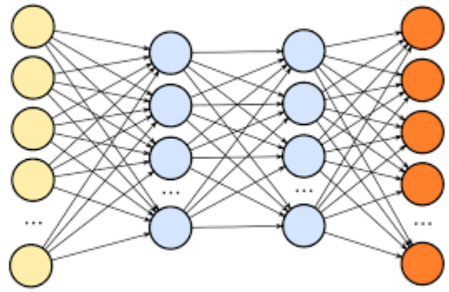

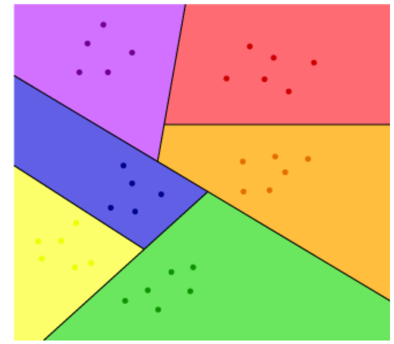

finding: young people are recommended pathways of lower effectiveness (high error rate)

fix the model!

of course, but maybe…

mental health info was missing for this population

go back to the data acquisition step, look for additional datasets

Mitigating urban homelessness

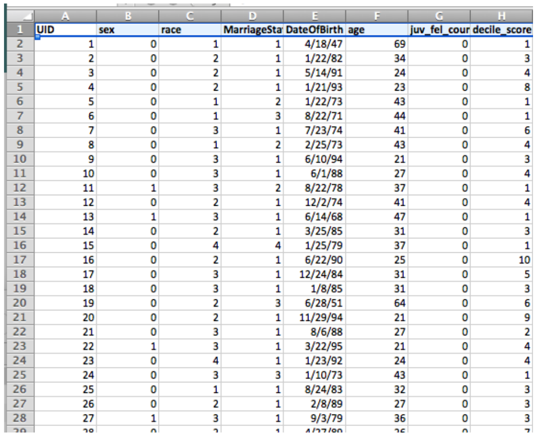

	A	B	C	D	E	F	G	H
1	UID	sex	race	MarriageSta	DateOfBirth	age	juv_fel_cour	decile_score
2	1	0	1	1	4/18/47	69	0	1
3	2	0	2	1	1/22/82	34	0	3
4	3	0	2	1	5/14/91	24	0	4
5	4	0	2	1	1/21/93	23	0	8
6	5	0	1	2	1/22/73	43	0	1
7	6	0	1	3	8/22/71	44	0	1
8	7	0	3	1	7/23/74	41	0	6
9	8	0	1	2	2/25/73	43	0	4
10	9	0	3	1	6/10/94	21	0	3
11	10	0	3	1	6/1/88	27	0	4
12	11	1	3	2	8/22/78	37	0	1
13	12	0	2	1	12/2/74	41	0	4
14	13	1	3	1	6/14/68	47	0	1
15	14	0	2	1	3/25/85	31	0	3
16	15	0	4	4	1/25/79	37	0	1
17	16	0	2	1	6/22/90	25	0	10
18	17	0	3	1	12/24/84	31	0	5
19	18	0	3	1	1/8/85	31	0	3
20	19	0	2	3	6/28/51	64	0	6
21	20	0	2	1	11/29/94	21	0	9
22	21	0	3	1	8/6/88	27	0	2
23	22	1	3	1	3/22/95	21	0	4
24	23	0	4	1	1/23/92	24	0	4
25	24	0	3	3	1/10/73	43	0	1
26	25	0	1	1	8/24/83	32	0	3
27	26	0	2	1	2/8/89	27	0	3
28	27	1	3	1	9/3/79	36	0	3

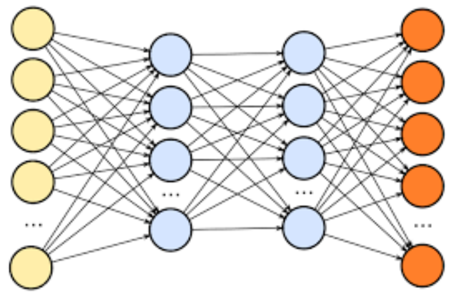

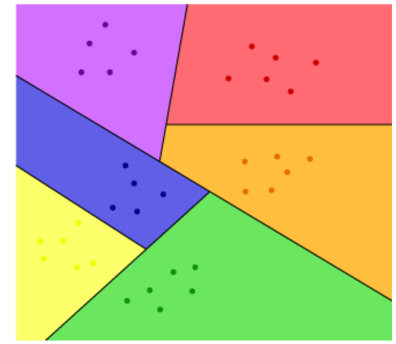

finding: minors are underrepresented in the input, compared to their actual proportion in the population (insufficient data)

unlikely to help!

fix the model??

minors data was not shared

go back to the data sharing step, help data providers share their data while adhering to laws and upholding the trust of the participants

interpretability: in the eye of the beholder

What are we explaining?

[J. Stoyanovich, J. Van Bavel, T. West; *NMI 2020*]

process (same for everyone? why is this the process?) vs. outcome

procedural justice aims to ensure that algorithms are perceived as fair and legitimate

data transparency is unique to algorithm-assisted decision-making, relates to the justification dimension of interpretability

To whom are we explaining and why?

[J. Stoyanovich, J. Van Bavel, T. West; *NMI 2020*]

accounting for the needs of different stakeholders

social identity - people trust their in-group members more

moral cognition - is a decision or outcome morally right or wrong?

How do we know that we explained well?

[J. Stoyanovich, J. Van Bavel, T. West; *NMI 2020*]

nutritional labels! :)

… but do they work?

To whom are we explaining and why?

[J. Stoyanovich, J. Van Bavel, T. West; *NMI 2020*]

accounting for the needs of different stakeholders

social identity - people trust their in-group members more

moral cognition - is a decision or outcome morally right or wrong?

back to decision-makers

Diversity in set selection

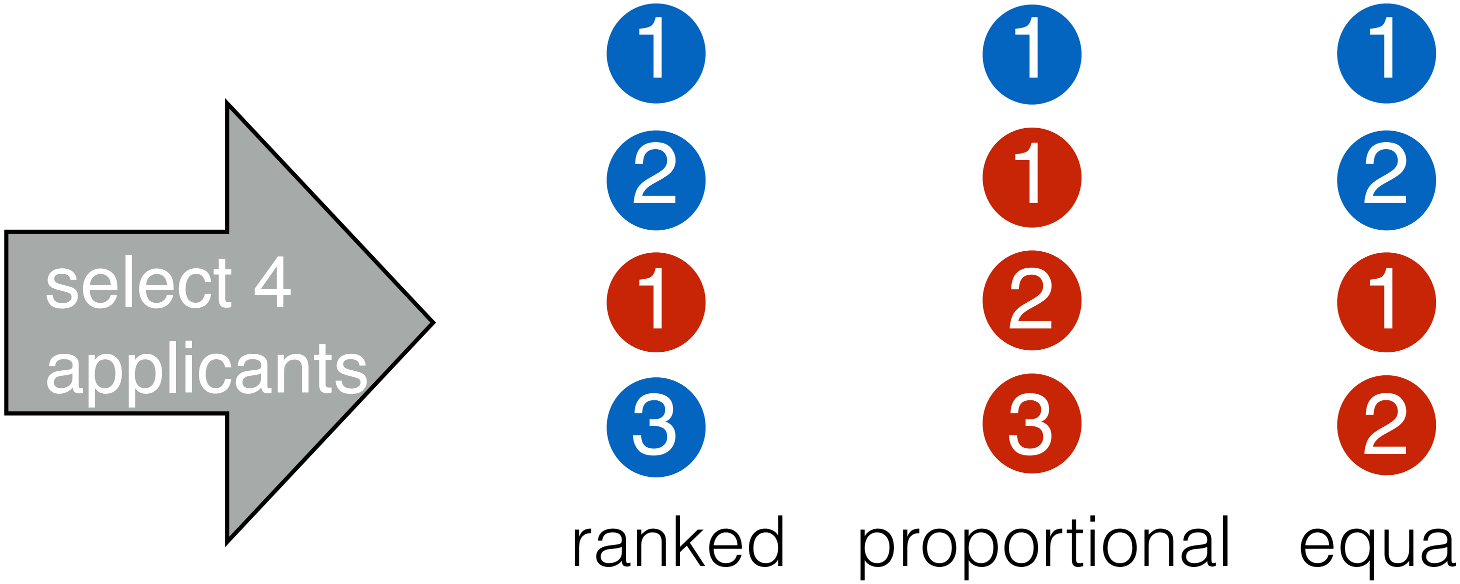

Can state all these as constraints:

for each category ***i***, pick K_i elements, with $floor_i \leq K_i \leq ceil_i$

[J. Stoyanovich, K. Yang, HV Jagadish; *EDBT 2018*]

Hiring a job candidate

Goal: Hire a candidate with a high score

4 1 3 2 5 7

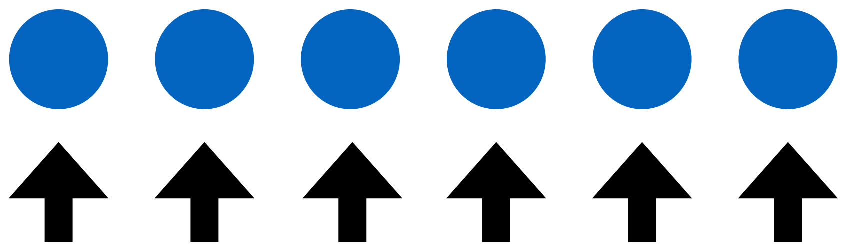

Candidates arrive one-by-one

A candidate's score is revealed when the candidate arrives

Decision to accept or reject a candidate made on the spot

The Secretary Problem

Goal: Design an algorithm for picking **one** element of a **randomly ordered** sequence, to maximize the probability of picking the **maximum element** of the entire sequence.

$$N = 6$$

$$S = \left\lfloor N/e \right\rfloor = 2$$

$$T = 4$$

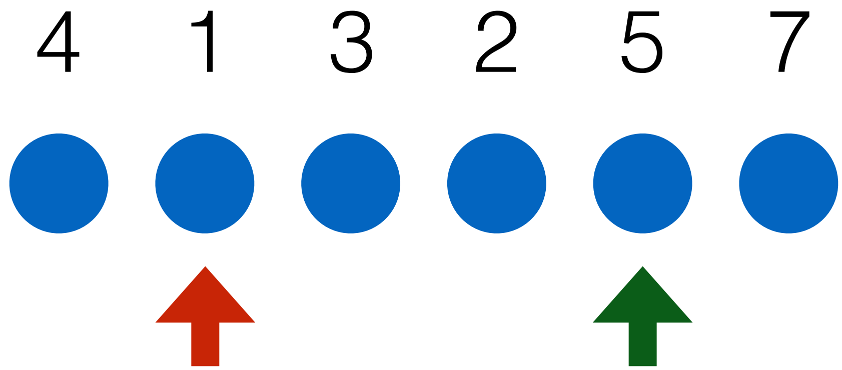

Competitive ratio

$$1/e$$

the best possible!

Consider, and reject, the first *S* candidates

Record *T*, the best seen score among the first *S* candidates

Accept the next candidate with score better than *T*

K-choice Secretary

[Babaioff et al., 2007]

Goal: Design an algorithm for picking **K** elements of a **randomly ordered** sequence, to maximize their **expected sum**.

$N = 6 \quad K = 2$

$S = \left\lfloor N/e \right\rfloor = 2$

$T = \{1, 4\}$

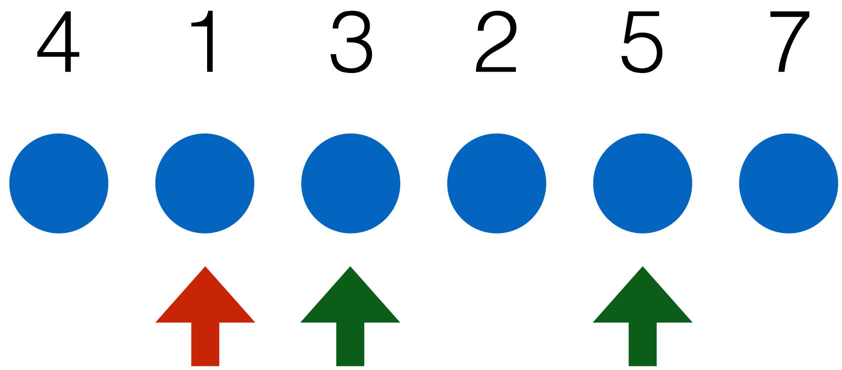

Competitive ratio

$1/e$

far from optimal

Consider, and reject, the first ***S*** candidates

Record ***K*** best scores among the first ***S*** candidates, call this ***T***

Whenever a candidate arrives whose score is higher than the minimum in ***T***, accept the candidate and delete the minimum from ***T***

Diverse K-choice Secretary

Goal: Design an algorithm for picking **K** elements of a **randomly ordered** sequence, to maximize their **expected sum**.

For each **category** ***i***, pick ***Kᵢ*** elements, with $floor_i \leq K_i \leq ceil_i$

6 4 1 8 2 3 1 2 9 5 7 5

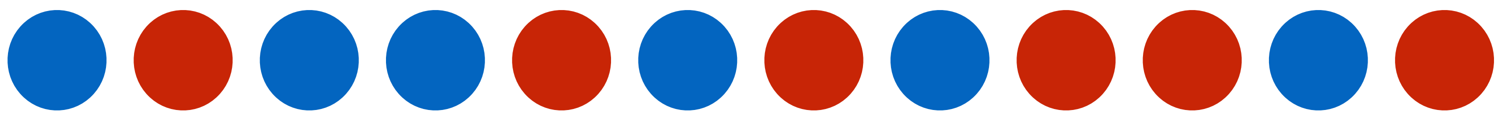

$$N_{red} = N_{blue} = 6$$

$$K = 3$$

$$1 \leq K_{red}, K_{blue} \leq 2$$

Accept *floor* items for each category from per-category streams

$$slack = K - (floor_{red} + floor_{blue})$$

Accept the remaining *slack* items irrespective of category membership, but subject to *ceil*

[J. Stoyanovich, K. Yang, HV Jagadish; *EDBT 2018*]

Diverse K-choice Secretary

$N_{red} = N_{blue} = 6$

$slack = 1$

$K = 3 \quad 1 \leq K_{red}, K_{blue} \leq 2$

$S_{red} = S_{blue} = 2 \quad S = 4$

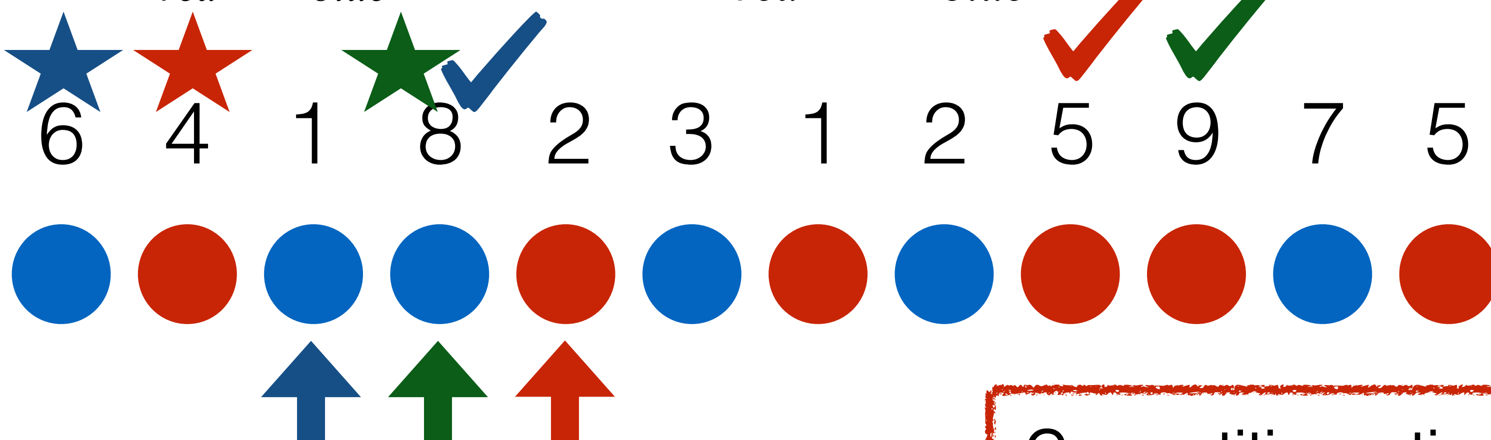

Competitive ratio

$1/e$

far from optimal

[J. Stoyanovich, K. Yang, HV Jagadish; *EDBT 2018*]

Per-category warm-up is crucial

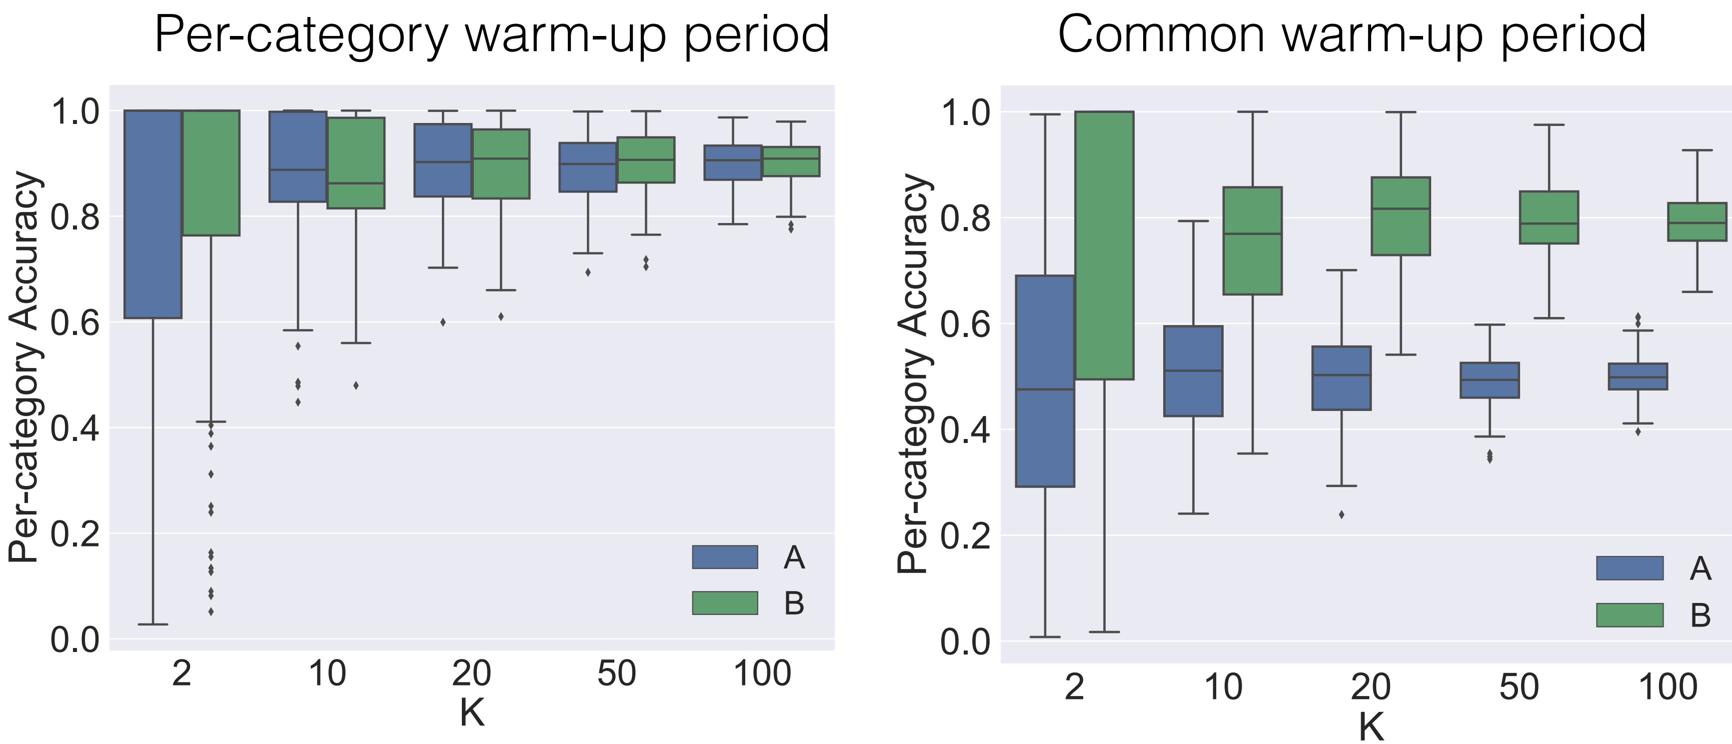

synthetic data with categories A and B, score depends on category, lower for A

diversity by design

[J. Stoyanovich, K. Yang, HV Jagadish; *EDBT 2018*]

Lack of diversity: harms and approaches

The New York Times

Artificial Intelligence's White Guy Problem

By KATE CRAWFORD JUNE 25, 2016

Like all technologies before it, artificial intelligence will reflect the values of its creators. So **inclusivity matters** — from who designs it to who sits on the company boards and which ethical perspectives are included.

Otherwise, **we risk constructing machine intelligence that mirrors a narrow and privileged vision of society**, with its old, familiar biases and stereotypes.

REVIEW

Diversity in Big Data: A Review

Marina Drosou[1] H.V. Jagadish[2] Evaggelia Pitoura[1] and Julia Stoyanovich[3,*]

Big Data
Volume 5 Number 2, 2017

DOI: 10.1089/big.2016.0054

Abstract

Big data technology offers unprecedented opportunities to society as a whole and also to its individual members. At the same time, this technology poses significant risks to those it overlooks. In this article, we give an overview of recent technical work on diversity, particularly in selection tasks, discuss connections between diversity and fairness, and identify promising directions for future work that will position diversity as an important component of a data-responsible society. We argue that diversity should come to the forefront of our discourse, for reasons that are both ethical—to mitigate the risks of exclusion—and utilitarian, to enable more powerful, accurate, and engaging data analysis and use.

Keywords: data; diversity; empirical studies; models and algorithms; responsibly

Score-based rankers

- tuple x in D; score(x): sum of attribute values, with non-negative weights (a common special case of monotone aggregation)
- weights subjectively chosen by a user: 0.5 g+ 0.5s, where g - normalized GPA, s - normalized SAT; why not 0.45 g + 0.55 s?

$\mathcal{D}$			f
id	x_1	x_2	$x_1 + x_2$
t_1	0.63	0.71	1.34
t_2	0.72	0.65	1.37
t_3	0.58	0.78	1.36
t_4	0.7	0.68	1.38
t_5	0.53	0.82	1.35
t_6	0.61	0.79	1.4

[A. Asudeh, HV Jagadish, G. Miklau, J. Stoyanovich; *VLDB 2019*]

Geometry of a (2D) ranker

$\mathcal{D}$			f
id	x_1	x_2	$x_1 + x_2$
t_1	0.63	0.71	1.34
t_2	0.72	0.65	1.37
t_3	0.58	0.78	1.36
t_4	0.7	0.68	1.38
t_5	0.53	0.82	1.35
t_6	0.61	0.79	1.4

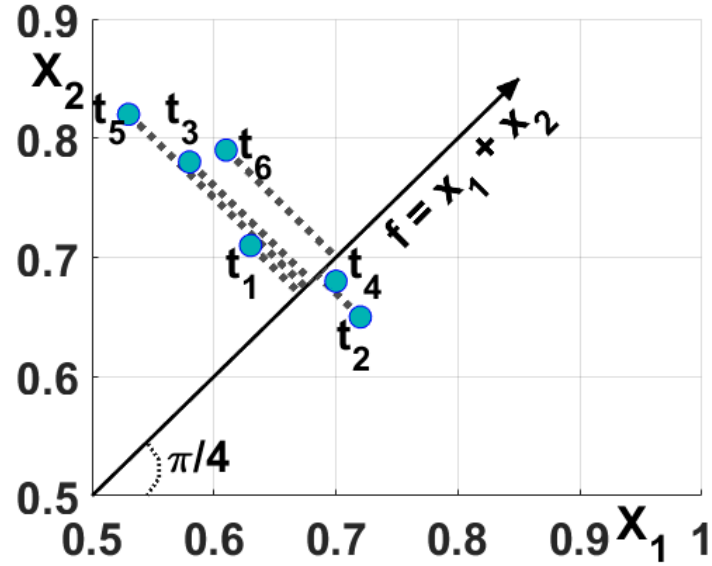

- tuples are points in 2D, scoring functions are rays starting from the origin
- to determine a ranking of the points, we read it off from the projections of the points onto the ray of the scoring function, walking the ray towards the origin
- examples: $f(x) = x_1 + x_2 \quad f(x) = x_1 \quad f(x) = x_2$

[A. Asudeh, HV Jagadish, G. Miklau, J. Stoyanovich; *VLDB 2019*]

Stability of a ranking

DEFINITION 2 (STABILITY OF $\mathfrak{r}$ AT $\mathcal{D}$). *Given a ranking $\mathfrak{r} \in \mathfrak{R}_{\mathcal{D}}$, the stability of $\mathfrak{r}$ is the proportion of ranking functions in $\mathcal{U}$ that generate $\mathfrak{r}$. That is, stability is the ratio of the volume of the ranking region of $\mathfrak{r}$ to the volume of $\mathcal{U}$. Formally:*

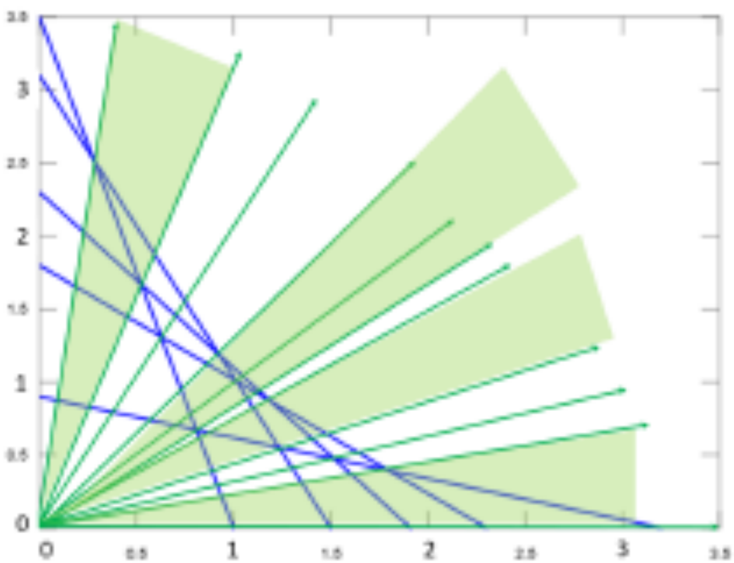

$$S_{\mathcal{D}}(\mathfrak{r}) = \frac{vol(R_{\mathcal{D}}(\mathfrak{r}))}{vol(\mathcal{U})}$$

most important finding:
FIFA rankings, used for seeding tournaments, are unstable! More in the paper.

[A. Asudeh, HV Jagadish, G. Miklau, J. Stoyanovich; *VLDB 2019*]

Ordering exchange

Key idea: only look at scoring functions that change the relative order between some pair of points. These are the only points where the oracle may change its mind!

$t_1 \langle 1,2 \rangle \quad t_2 \langle 2,1 \rangle \qquad t_2 \succ_x t_1 \qquad \boxed{t_2 =_{x+y} t_1} \qquad t_2 \prec_y t_1$

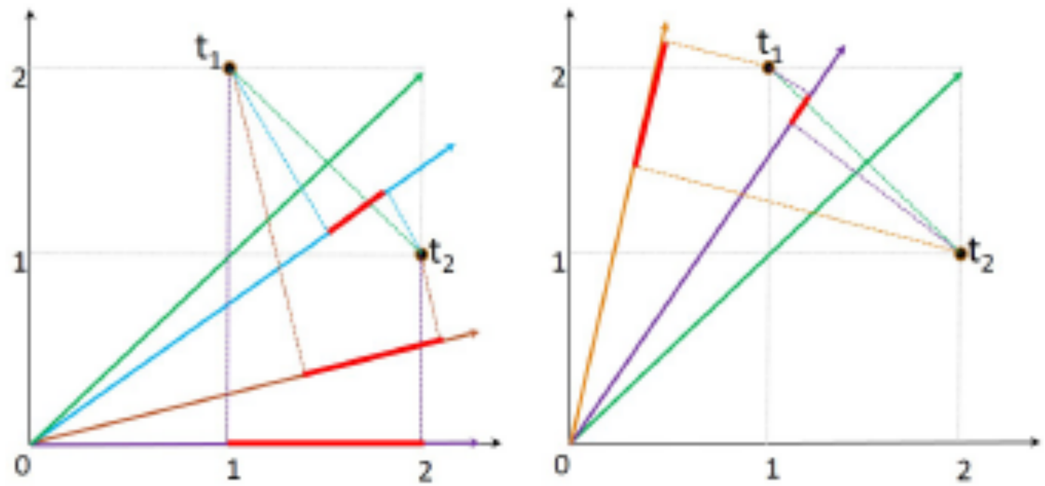

An **ordering exchange** is a set of functions that score a pair of points equally. In 2D, it corresponds to a single function.

[A. Asudeh, HV Jagadish, G. Miklau, J. Stoyanovich; *VLDB 2019*]

Working with the geometry

Step 1: pre-compute an index over the space

Step 2: efficiently answer questions at query time

- Give a list of stable regions
- Give stable regions closest to my scoring function
- Interrogate stability or fairness of a scoring function
- (Similar methods to compute fair or diverse regions)

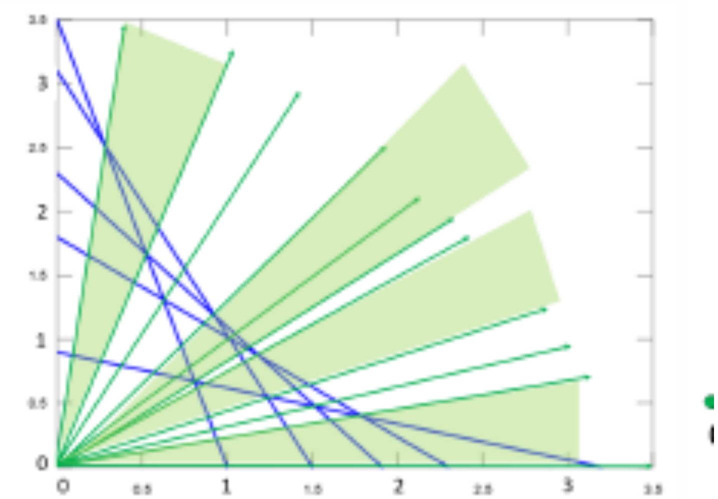

[A. Asudeh, HV Jagadish, G. Miklau, J. Stoyanovich; *VLDB 2019*]

taking responsibility

Personal responsibility?

NATURE | NEWS

nature International weekly journal of science

Home | News & Comment | Research | Careers & Jobs | Current Issue | Archive | Audio & Video | For Authors

Archive > Volume 515 > Issue 7526 > News > Article

Italian seismologists cleared of manslaughter

Appeals court says six scientists did not cause deaths in 2009 L'Aquila earthquake and cuts sentence of a government official.

Alison Abbott & Nicola Nosengo

10 November 2014

Six seismologists accused of misleading the public about the risk of an earthquake in Italy were cleared of manslaughter on 10 November. An appeals court overturned their six-year prison sentences and reduced to two years the sentence for a government official who had been convicted with them.

The magnitude-6.3 earthquake struck the historic town of L'Aquila in the early hours of 6 April 2009, killing more than 300 people.

Personal responsibility?

NATURE | NEWS

Italian seismologists cleared of manslaughter

Appeals court says six scientists did not cause deaths in 2009 L'Aquila earthquake and cuts sentence of a government official.

Alison Abbott & **Nicola Nosengo**

10 November 2014

The finding by a three-judge appeals court prompted many L'Aquila citizens who were waiting outside the courtroom to react with rage, shouting "shame" and saying that the Italian state had just acquitted itself, local media reported. But it **comes as a relief to scientists around the world who had been following the unprecedented case with alarm.**

"We don't want to have to be worried about the possibility of being prosecuted if we give advice on earthquakes," says seismologist Ian Main of the University of Edinburgh, UK. "That would discourage giving honest opinion."