Fair Ranking with Biased Data

Ashudeep Singh, Himank Yadav, Zhengxiao Du, Magd Bayoumi, Yi Su,
Thorsten Joachims

Department of Computer Science, Cornell University
Department of Information Science, Cornell University

Funded in part through NSF Awards IIS-2008139, IIS-1615706, IIS-1901168.

Ranking in Online Systems

Ranking function π that ranks items for context x.

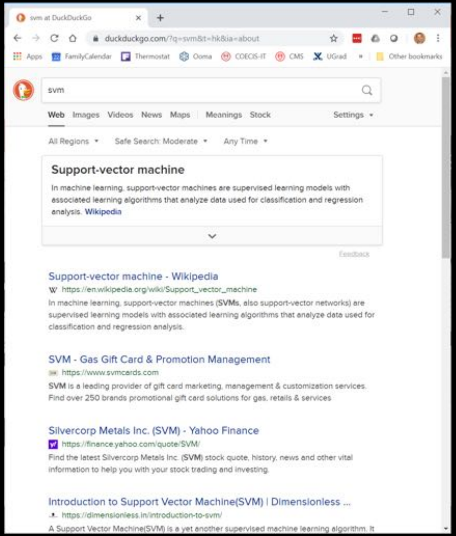

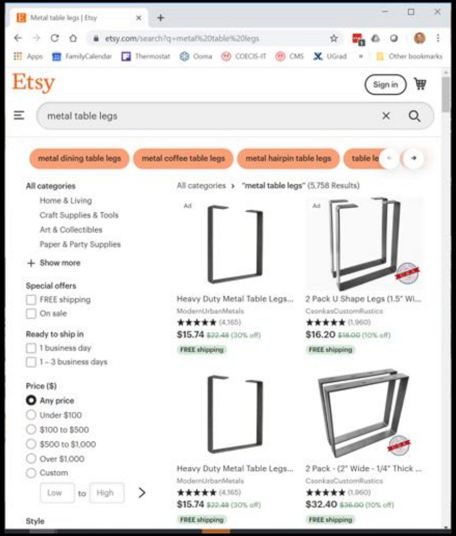

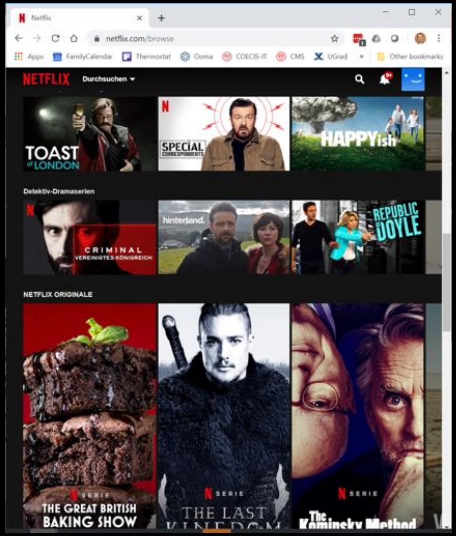

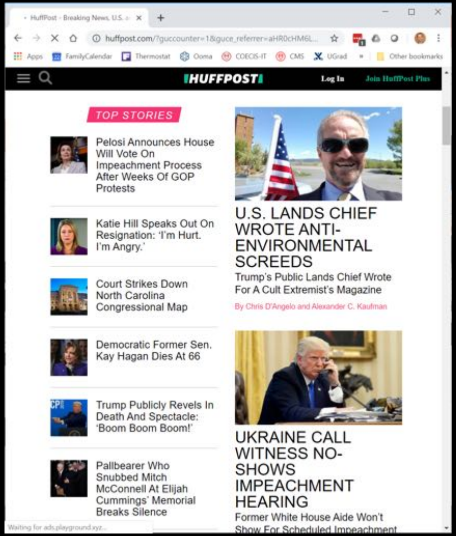

What is the ideal ranking?

Goal: Maximize utility of rankings to the users.

©FastCompany

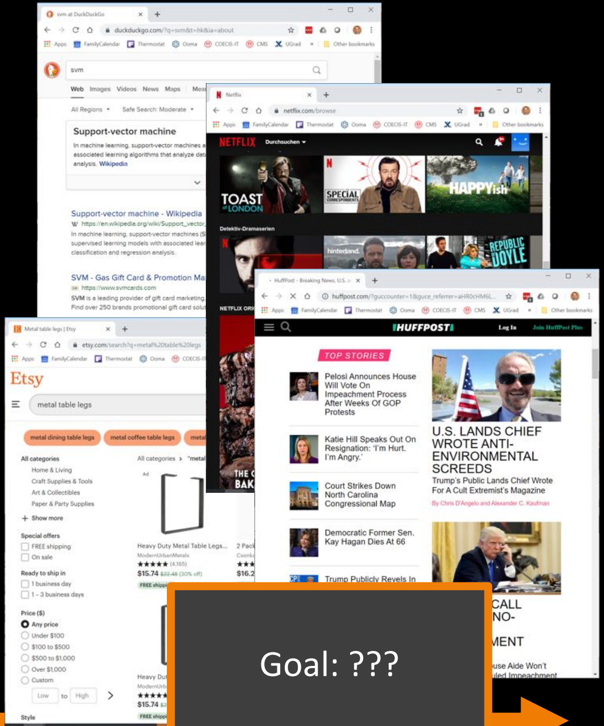

Goal: ???

1960

1994

2020

Two-Sided Market

Online Retail

- Utility to Users:

  Customers find products they want

- Utility to Items:

  Sellers get revenue

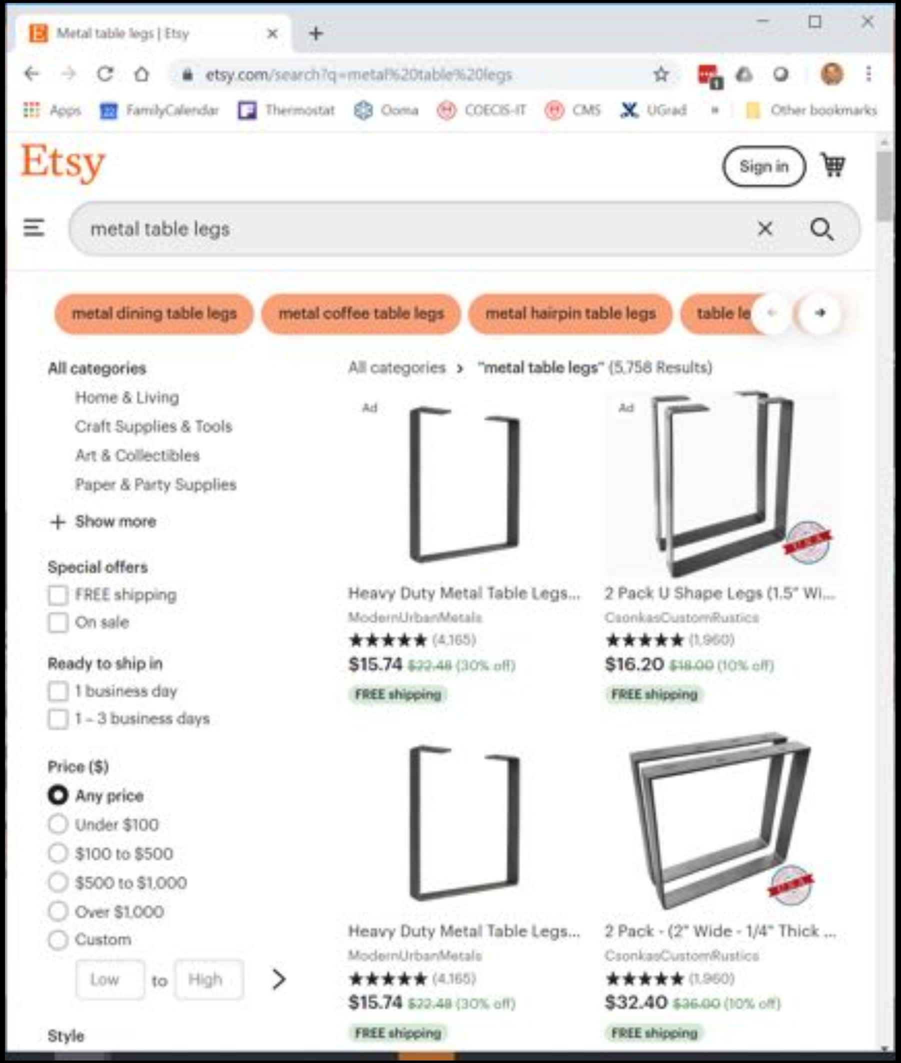

Two-Sided Market

Music Streaming

- Utility to Users:

  Customers find music they enjoy

- Utility to Items:

  Artists get streaming revenue

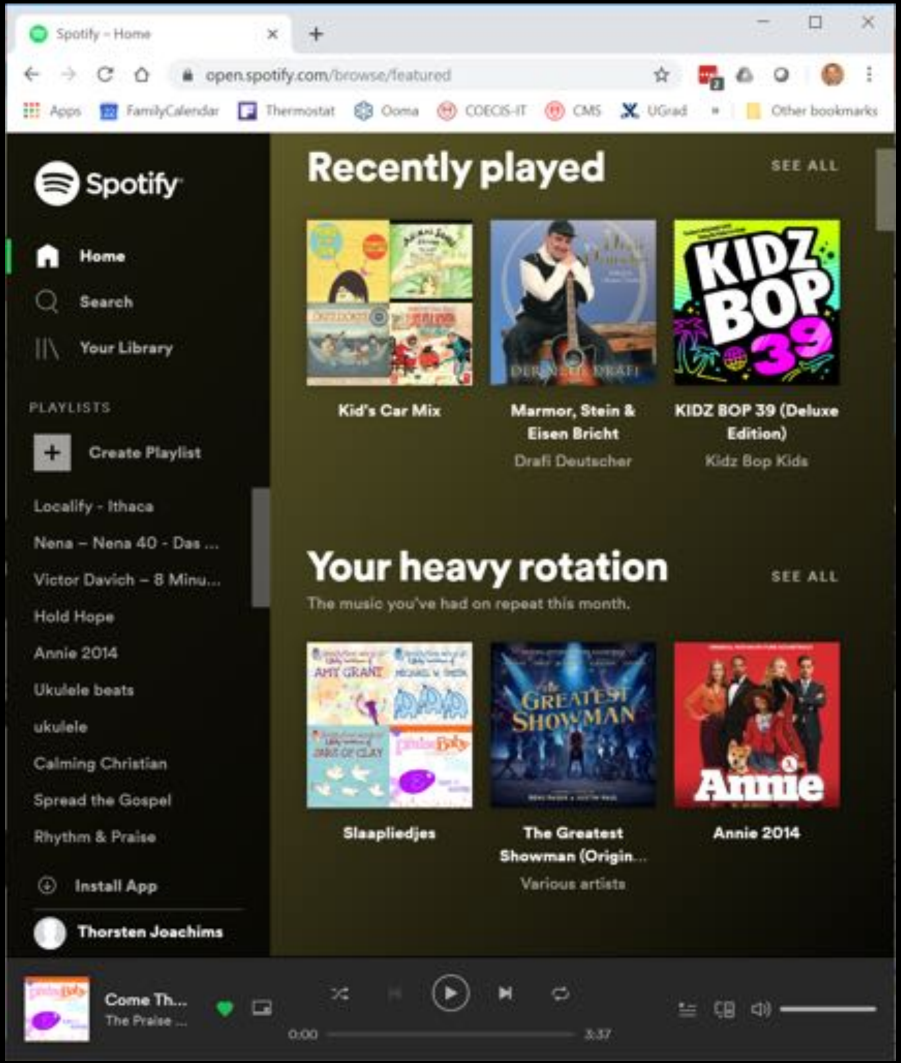

Two-Sided Market

Research Papers

- Utility to Users:

  Readers find relevant articles

- Utility to Items:

  Writers get their voice out (and tenure)

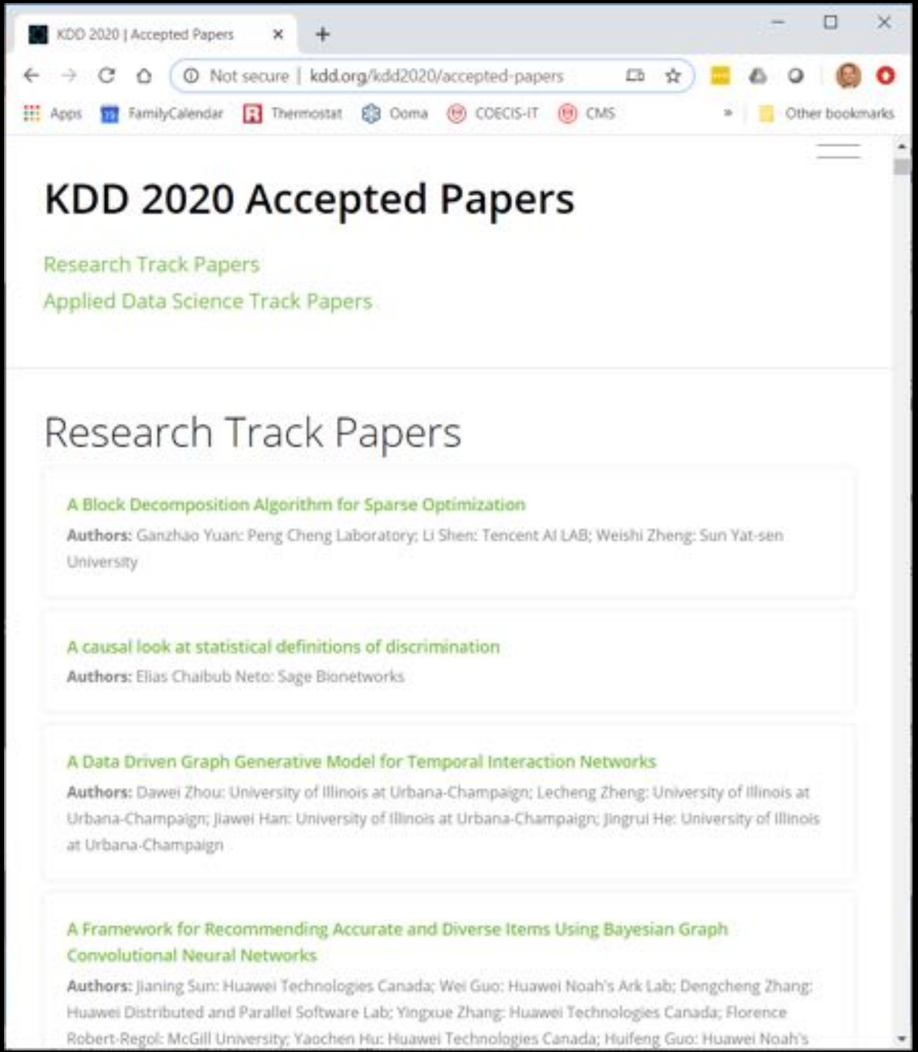

Maximizing Utility to Users

Probability Ranking Principle [Robertson, 1977]:

- Rank documents by probability of relevance → y^*
- For virtually any measure U of ranking quality

$$y^* \coloneqq \operatorname{argmax}_y [\mathrm{U}(y|x)]$$

Dynamics of Utility Maximization

Probability Ranking Principle:

- Rank documents by probability of relevance → y^* [Robertson, 1977]
- For virtually any measure U of ranking quality

$$y^* := \mathrm{argmax}_y[\mathrm{U}(y|x)]$$

- Are rankings fair/desirable?

Top News Stories		
Rank	Item	P(read)
1	Times 1	50.99
2	Times 2	50.98
3	Times 3	50.97
…	…	…
100	Review 1	49.99
101	Review 2	49.98
102	Review 3	49.97
…	…	…

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- be fair to the items
- satisfy legal requirements
- shape market dynamics (e.g. superstar economics, spam, polarization)

Exogenous Factors

How to estimate merit without biases like

- position bias
- trust bias
- uncertainty bias
- stereotypes

Position-Based Exposure Model

Definition:

Exposure e_j is the probability a users observes the item at position j.

$$Exp(G|x,y) = \sum_{j \in G} e_j$$

How to estimate?

- Eye tracking [Joachims et al. 2007]
- Intervention studies [Joachims et al. 2017]
- Intervention harvesting [Agarwal et al. 2019] [Fang et al. 2019]

Rank	Exposure P(observe)
1	e_1
2	e_2
3	e_3
…	…
100	e_{100}
101	e_{101}
102	e_{102}
…	…

Fairness Disparity

Goal: $Exp(G|x,y) = f\big(Rel(G|x)\big)$

Example: Make exposure proportional to relevance (per group)

$$\frac{Exp(G_0|x,y)}{Exp(G_1|x,y)} = \frac{Rel(G_0|x)}{Rel(G_1|x)}$$

Disparity: $D(y|x) = \left|Exp(G|x,y) - f\big(Rel(G|x)\big)\right|$

[Singh & Joachims, 2018] [Biega et al., 2018]

Learning Fair Ranking Policies

Goal: Policy π that maximizes expected utility U with small disparity D.

$$\pi^* = \text{argmax}_\pi E_x[U(\pi|x)] \quad s.t.\ E_x[D(\pi|x)] \leq \delta$$

Learning: Empirical Risk Minimization

$$\hat{\pi} = \text{argmax}_\pi \frac{1}{n}\sum_{i=1}^{n} U(\pi|x_i) \quad s.t. \frac{1}{n}\sum_{i=1}^{n} D(\pi|x_i) \leq \delta$$

→ Lagrange multiplier

$$\hat{\pi} = \text{argmax}_\pi \frac{1}{n}\sum_{i=1}^{n} U(\pi|x_i) - \lambda\frac{1}{n}\sum_{i=1}^{n} D(\pi|x_i)$$

[Singh & Joachims, 2019]

Stochastic Ranking Policies

- Policy:
  $\pi(y|x)$ is conditional distribution over rankings.
- Utility:
  $$U(\pi|x) = \sum_y \mathrm{U}(y|x)\pi(y|x)$$
- Exposure:
  $$\mathrm{Exp}(G|x,\pi) = \sum_{j \in G} \sum_{y} e_{rank(j|y)} \pi(y|x)$$

π

y_1	y_2	y_3	y_4
A	B	A	B
B	A	C	C
C	C	B	A
D	D	D	G
E	E	E	F
F	F	F	E
G	G	G	D
0.52	0.23	0.20	0.05

[Singh & Joachims, 2018]

Policy Training

Training objective:

$$\hat{\pi} = \operatorname{argmax}_{\pi} \frac{1}{n}\sum_{i=1}^{n} U(\pi|x_i) - \lambda \frac{1}{n}\sum_{i=1}^{n} D(\pi|x_i)$$

Policy class:

- Plackett-Luce $\pi_w(y|x) = PL(s_1, \dots, s_k)$ with per-item scoring model $\mathrm{s_j} = s\left(y_j \middle| x, w\right)$

Training algorithm:

- Policy gradient with Monte-Carlo estimates of gradient.
- Entropy regularization.
- Variance reduction.

[Singh & Joachims, 2019]

Experiment

Data
- Yahoo LTR Challenge

Fairness
- Proportional exposure
- Individual fairness

Ranking policy
- Plackett-Luce
- Deep network scorer

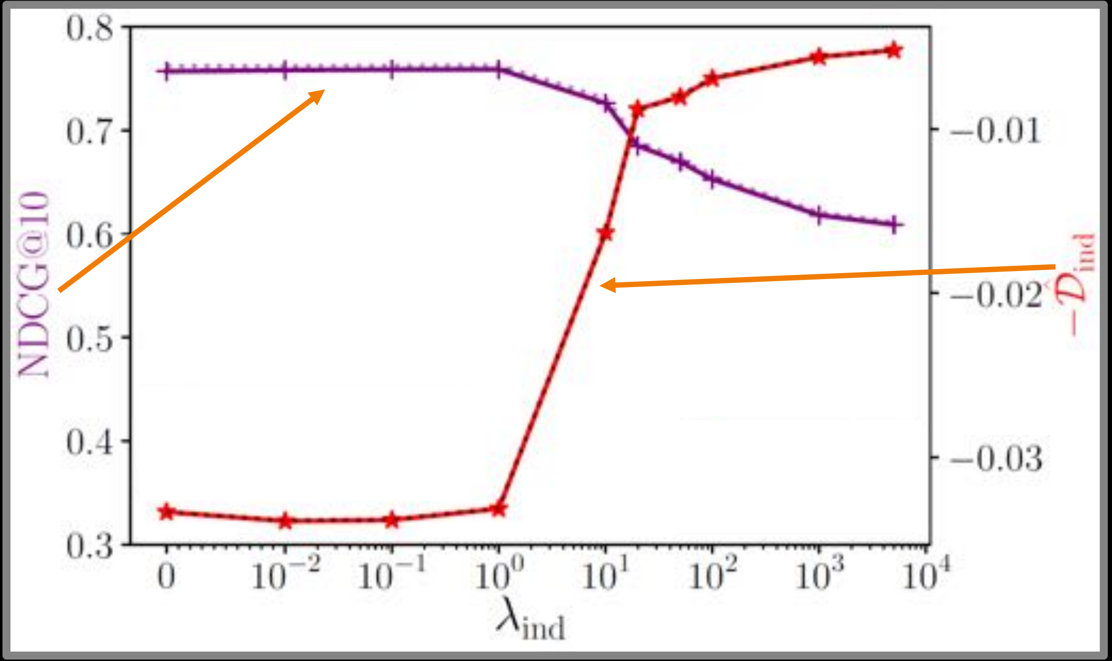

→ Generalizes to be fair on test data.

[Singh & Joachims, 2019]

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- be fair to the items
- satisfy legal requirements
- shape market dynamics (e.g. superstar economics, spam, polarization)

Exogenous Factors

How to estimate merit without biases like

- position bias
- trust bias
- uncertainty bias
- stereotypes

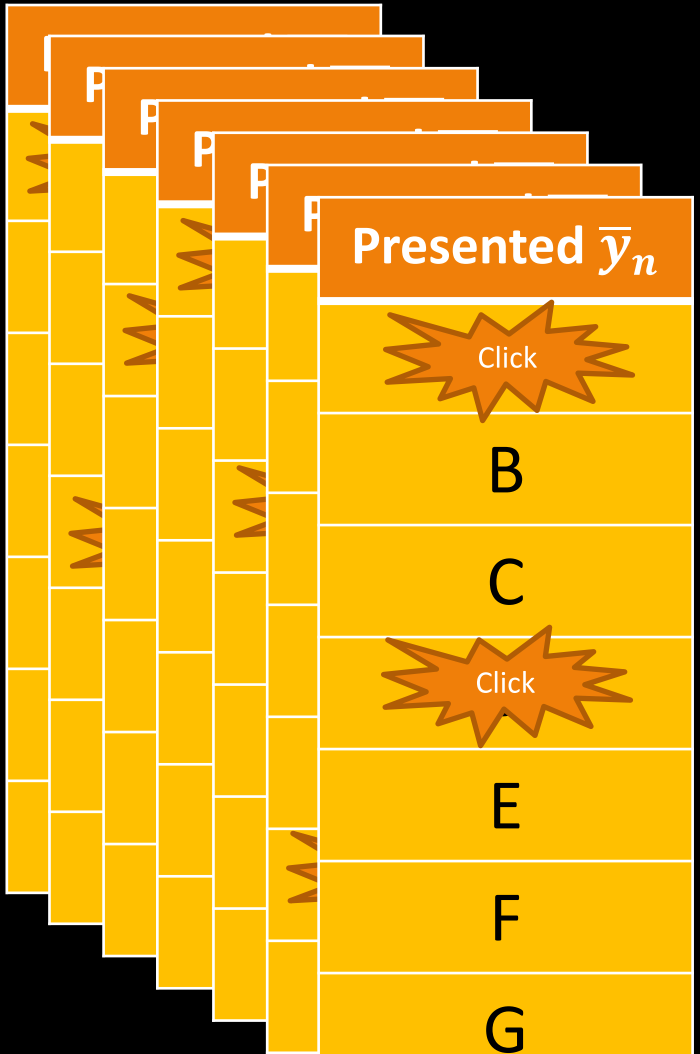

Interaction Feedback

Data

- Query distribution: $x_j \sim \boldsymbol{P(X)}$
- Deployed ranker: $\bar{y}_j \sim \boldsymbol{\pi_0}(y|x_j)$
- Feedback: clicks, purchases, plays, reads

→ Feedback is biased!

Modeling Position Bias

- Assume:
  - Click implies observed and relevant:

$$(click_i = 1) \leftrightarrow (obs_i = 1) \wedge (rel_i = 1)$$

- Problem:
  - No click can mean not relevant OR not observed

$$(click_i = 0) \leftrightarrow (obs_i = 0) \vee (rel_i = 0)$$

→ Understand observation mechanism

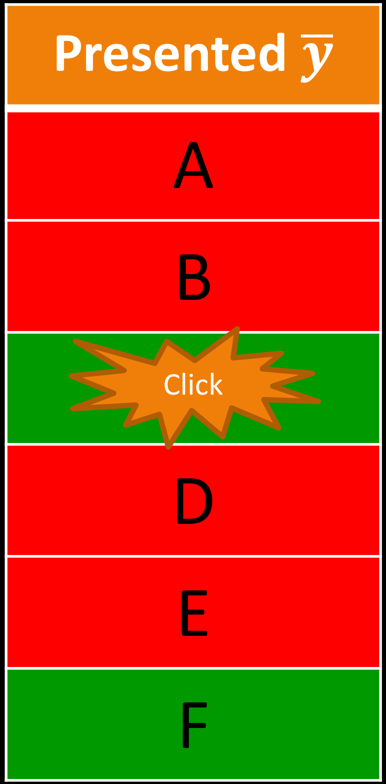

Inverse Propensity Score Estimators

- Observation Propensities
  - $Q(obs_j = 1|x, \bar{y})$
  - Random variable $obs_j \in \{0,1\}$ indicates whether relevance label rel_j is observed.
  - Can use position-based exposure $Q(obs_j = 1|x, \bar{y}) = e_j$
- Inverse Propensity Score (IPS) Weighting
  - Utility: $\hat{\mathrm{U}}(y|x) = \sum_j g(rank(j|y)) \frac{click(j|x)}{e_j}$ (e.g. DCG)
  - Relevance: $\widehat{Rel}(G|x) = \sum_{j \in G} \frac{click(j|x)}{e_j}$

→ Unbiased!
In expectation independent of past rankings.

Presented $\bar{y}$	Q
A	1.0
B	0.8
C	0.5
D	0.2
E	0.2
F	0.2

[Joachims et al., 2017] [Yadav et al., 2020]

Fair Policy Training

Training objective:

$$\hat{\pi} = \operatorname{argmax}_{\pi} \frac{1}{n}\sum_{i=1}^{n} \hat{U}(\pi|x_i) - \lambda \frac{1}{n}\sum_{i=1}^{n} \hat{D}(\pi|x_i)$$

Utility

– Unbiased $\hat{\mathrm{U}}(y|x)$ gives unbiased $\hat{U}(\pi|x_i)$

Disparity

– Average relevance $\widehat{Rel}(G) = \sum_x \widehat{Rel}(G|x)$

– Amortized group disparity (similar to [Biega et al., 2018])

$$\hat{D}(y|x) = \widehat{Rel}(G_1)Exp(G_0|x) - \widehat{Rel}(G_0)Exp(G_1|x)$$

[Yadav et al., 2020]

Experiment

Data
- Microsoft LTR Corpus

Fairness
- Amortized proportional exposure
- Group fairness

Ranking policy
- Plackett-Luce
- Linear scorer

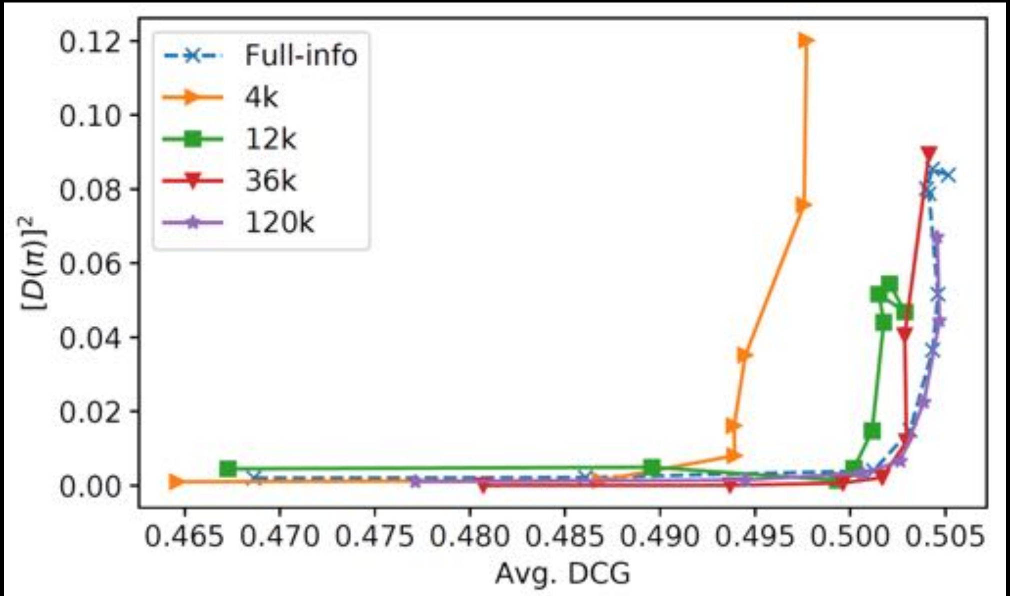

[Yadav et al., 2020]

Comparison

- Group blind
  - Fairness through unawareness
- Post processing
  - IPS regression
  - Biega et al. fairness
- Fair-PG-Rank
  - Method from before

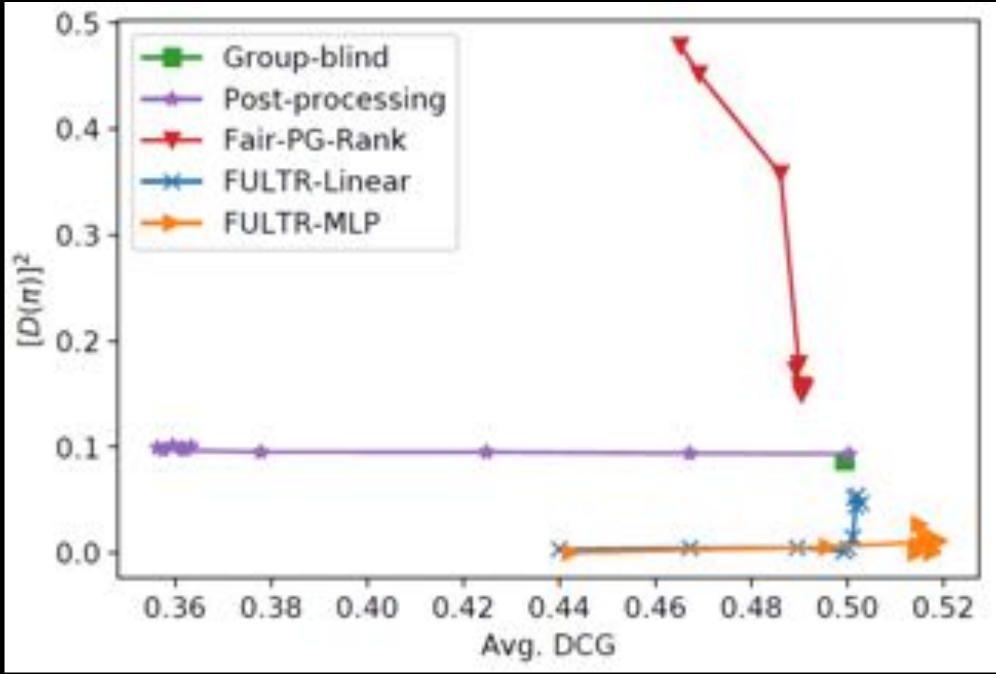

[Yadav et al., 2020]

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- be fair to the items
- satisfy legal requirements
- shape market dynamics (e.g. superstar economics, spam, polarization)

Exogenous Factors

How to estimate merit without biases like

- position bias
- trust bias
- uncertainty bias
- stereotypes

Matching Markets

Employer	Preference
Z	A > D > …
Y	C > A > …
X	E > C > …
W	A > B > …
V	A > D > …
…	…

←

Job Recommender

→

Applicant	Preference
A	X > Z > …
B	W > V > …
C	Y > X > …
D	Y > Z > …
E	V > Z > …
…	…

→ Multi-sided Preferences, Fairness, and Social Welfare.

[Tu et al., 2014] [Hopcroft et al., 2011]

Simulation Experiment

Effect on Market

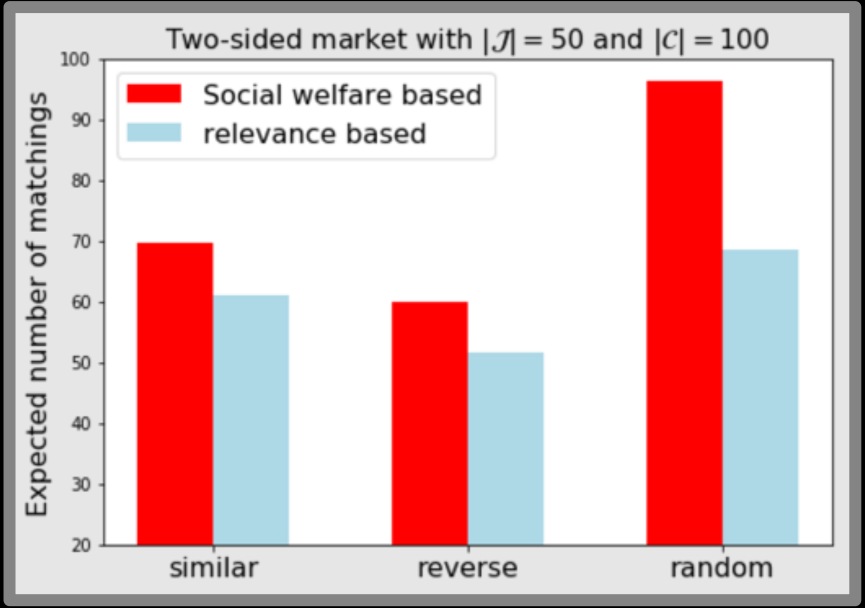

Effect on Individuals

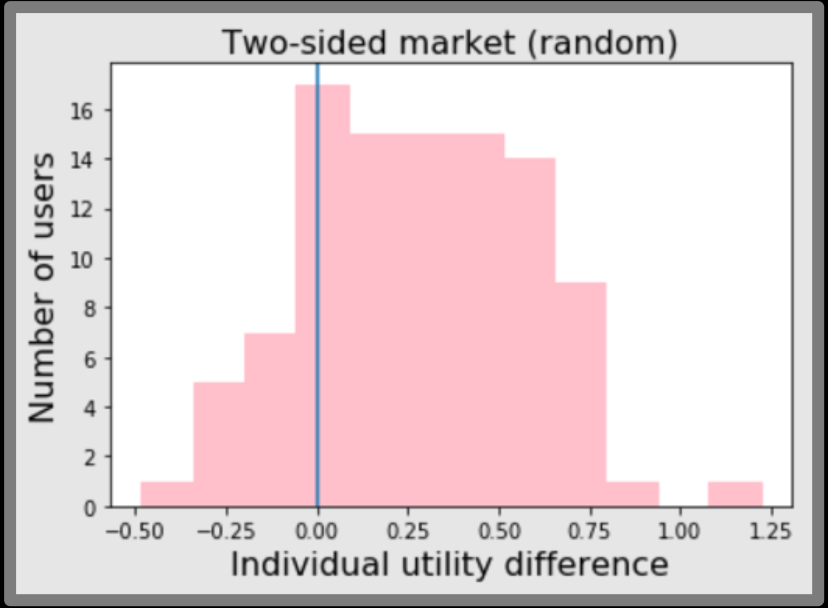

[Su et al., 2021]

Research Agenda for Ranking

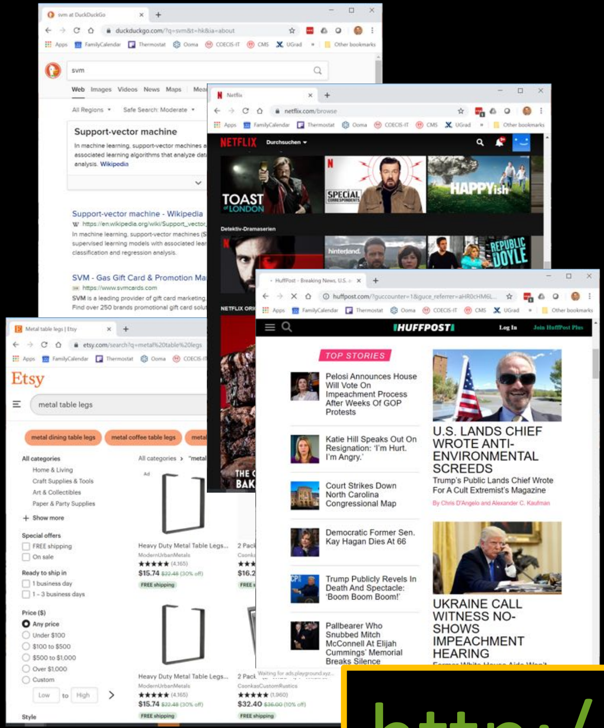

- Fairness to items
- Fairness to user groups
- Market-level objectives
- Long-term dynamics
- Transparency
- Privacy

http://www.joachims.org