LinkedIn

A Framework for Fairness in Two-Sided Marketplaces

Kinjal Basu, Cyrus DiCiccio, Heloise Logan, Noureddine El Karoui
LinkedIn Corporation

Presented by:
Omid Memarrast

Outline

- Introduction
- Background
- Two-Sided Fair Marketplace
- Experiments
- Conclusion

Motivation

Suggestions: **one gender!**

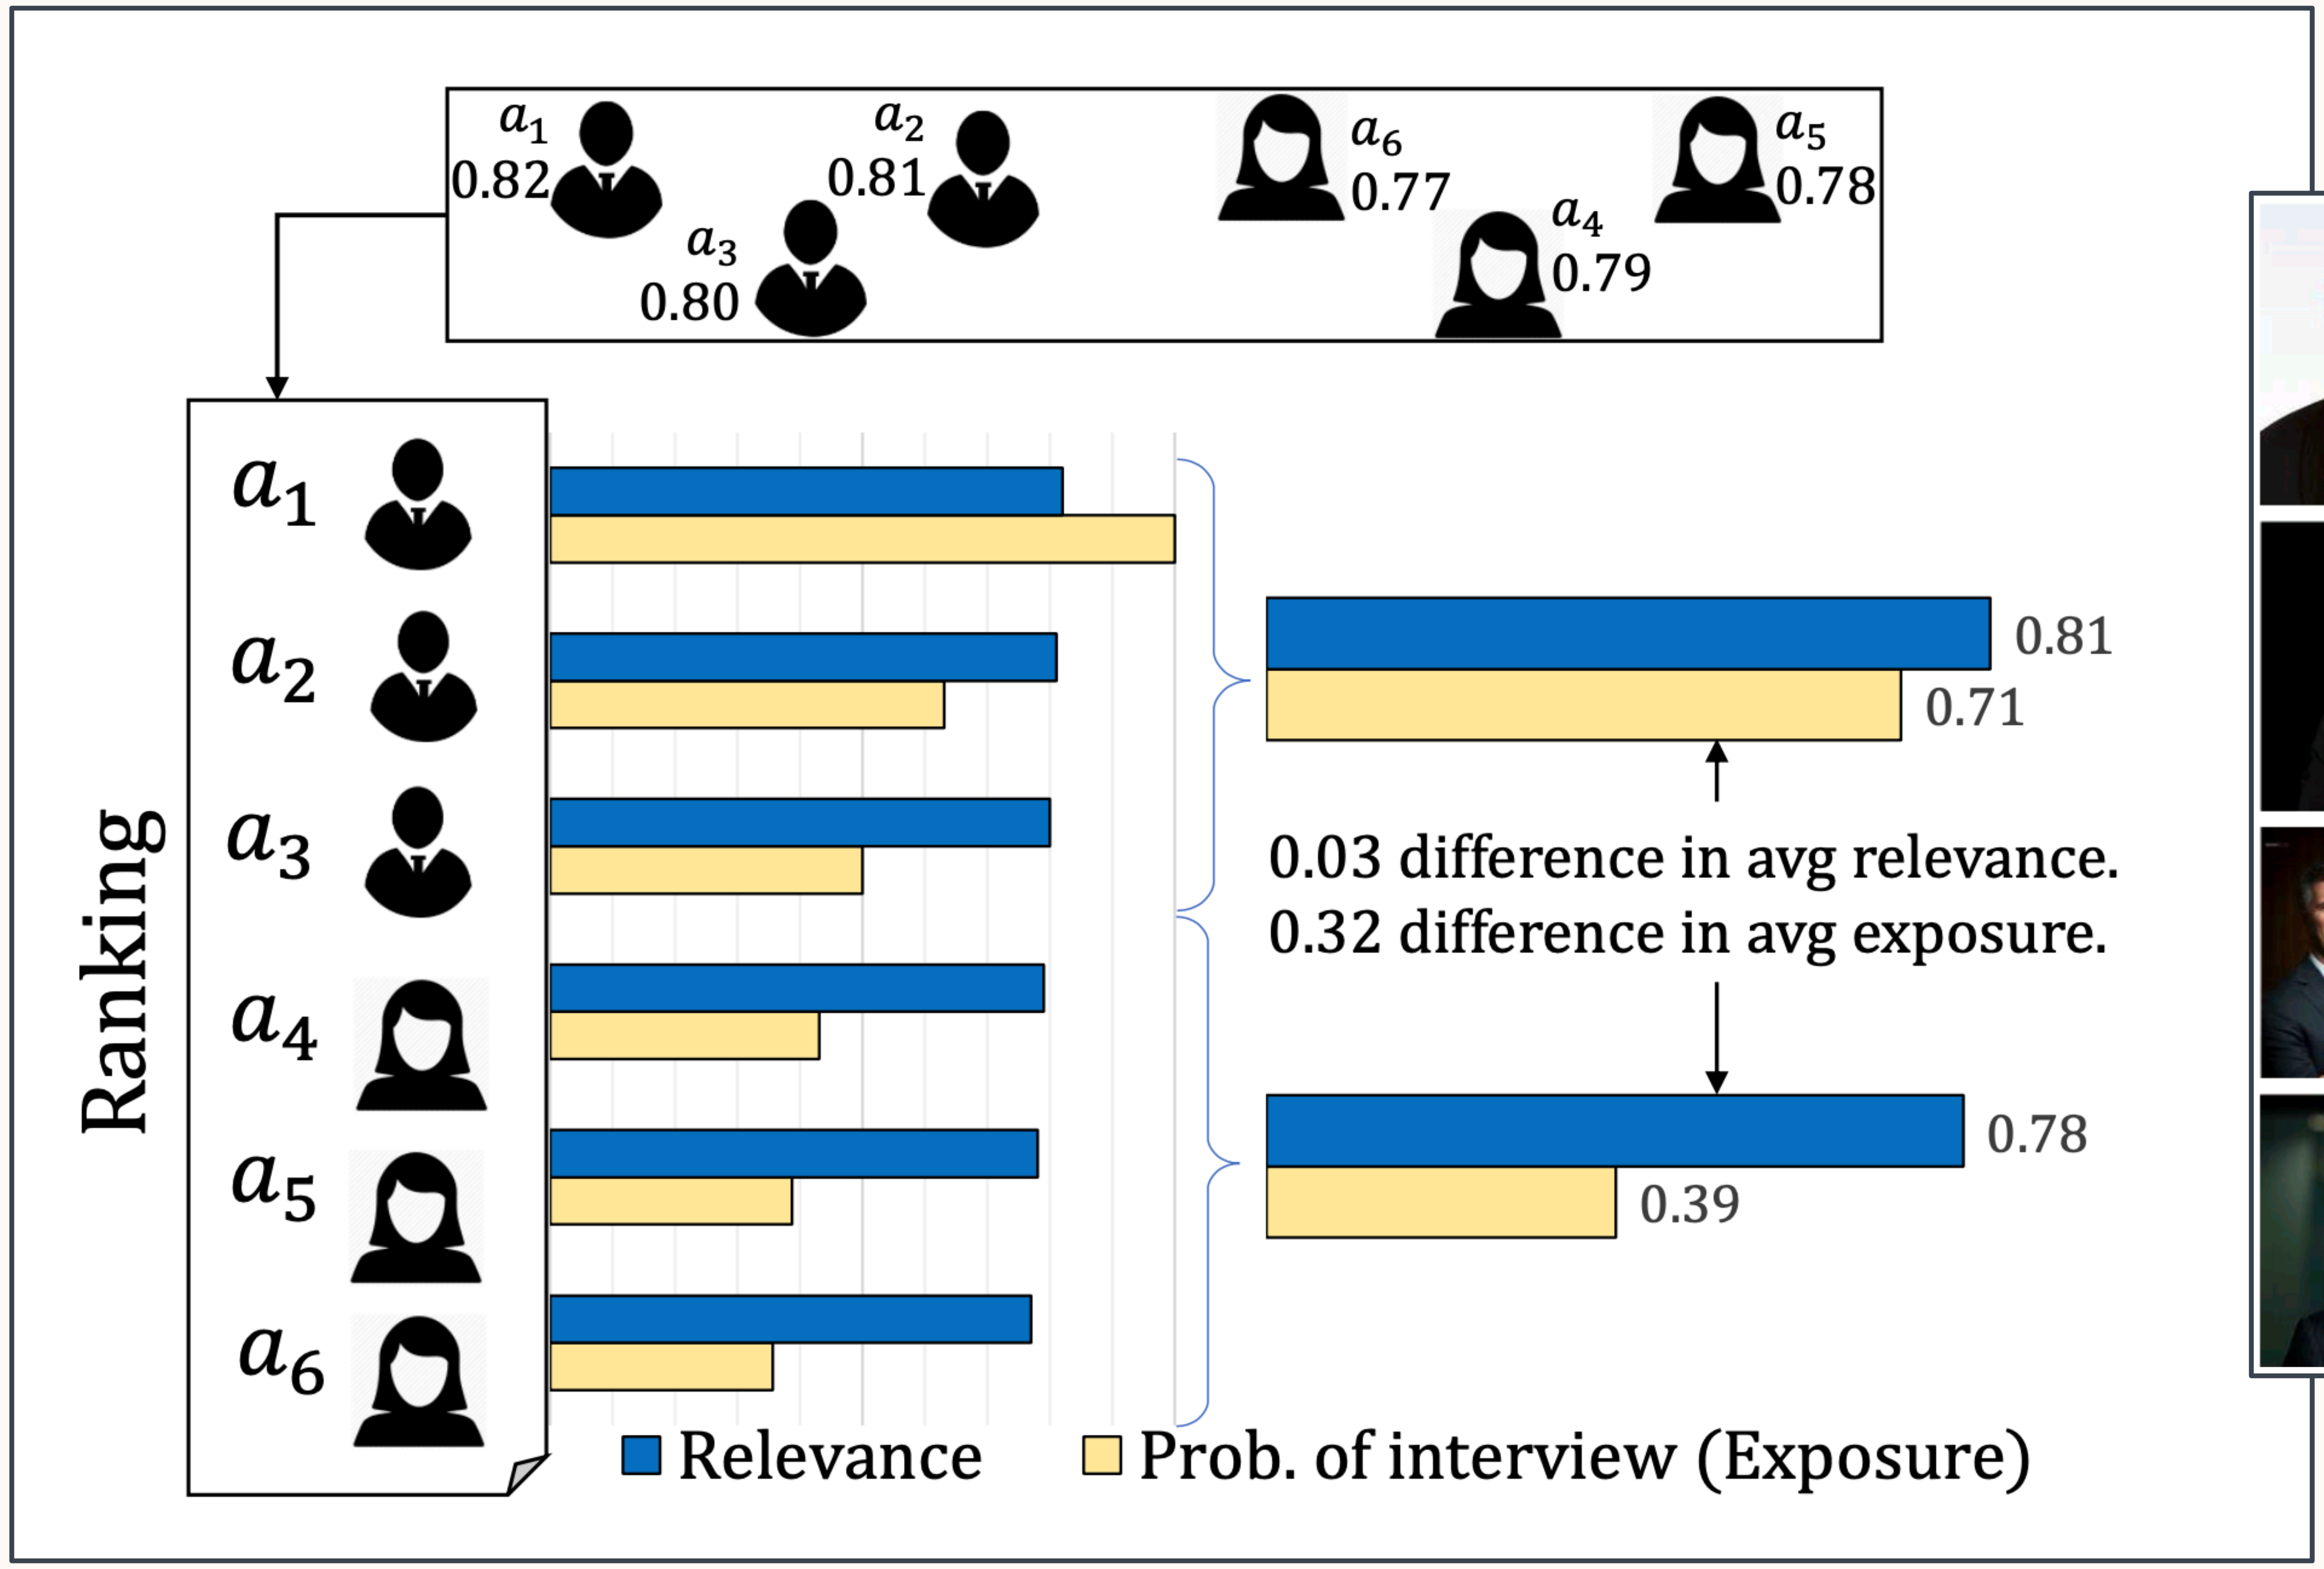

Job Seeker Example

Search for CEO in a serach engine

Singh, Ashudeep, and Thorsten Joachims. "Fairness of exposure in rankings." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2018.

Introduction

Recommender systems

Single-sided

members viewing:
- Items
- Products
- Jobs
- Movies
- Restaurants

Two-sided member-to-member marketplace

when members can serve multiple functions:

- feed ranking
- people or friend recommendations
- search systems
- recruiters searching for job candidates

Introduction

Goal:

Guarantee that the rankings are fair to:

- The *source* members initiating the queries
- The *destination* members who are being returned by the query

How?

- Define a multi-session utility and add fairness constraints for *destination* members across multiple sessions
- Add fairness constraints for *source* members
- Scale for large-scale recommender systems by using a duality idea

People You May Know (Linkedin)

Suggestions: **one gender!**

- PYMK is a Member-Member Marketplace
- Fairness based on Gender attribute (Similar setting for IM, FM)
- Destination Side Fairness
  - Directly observable by members
- Source Side Fairness
  - Unobservable by individual members
  - Can be exposed if audited

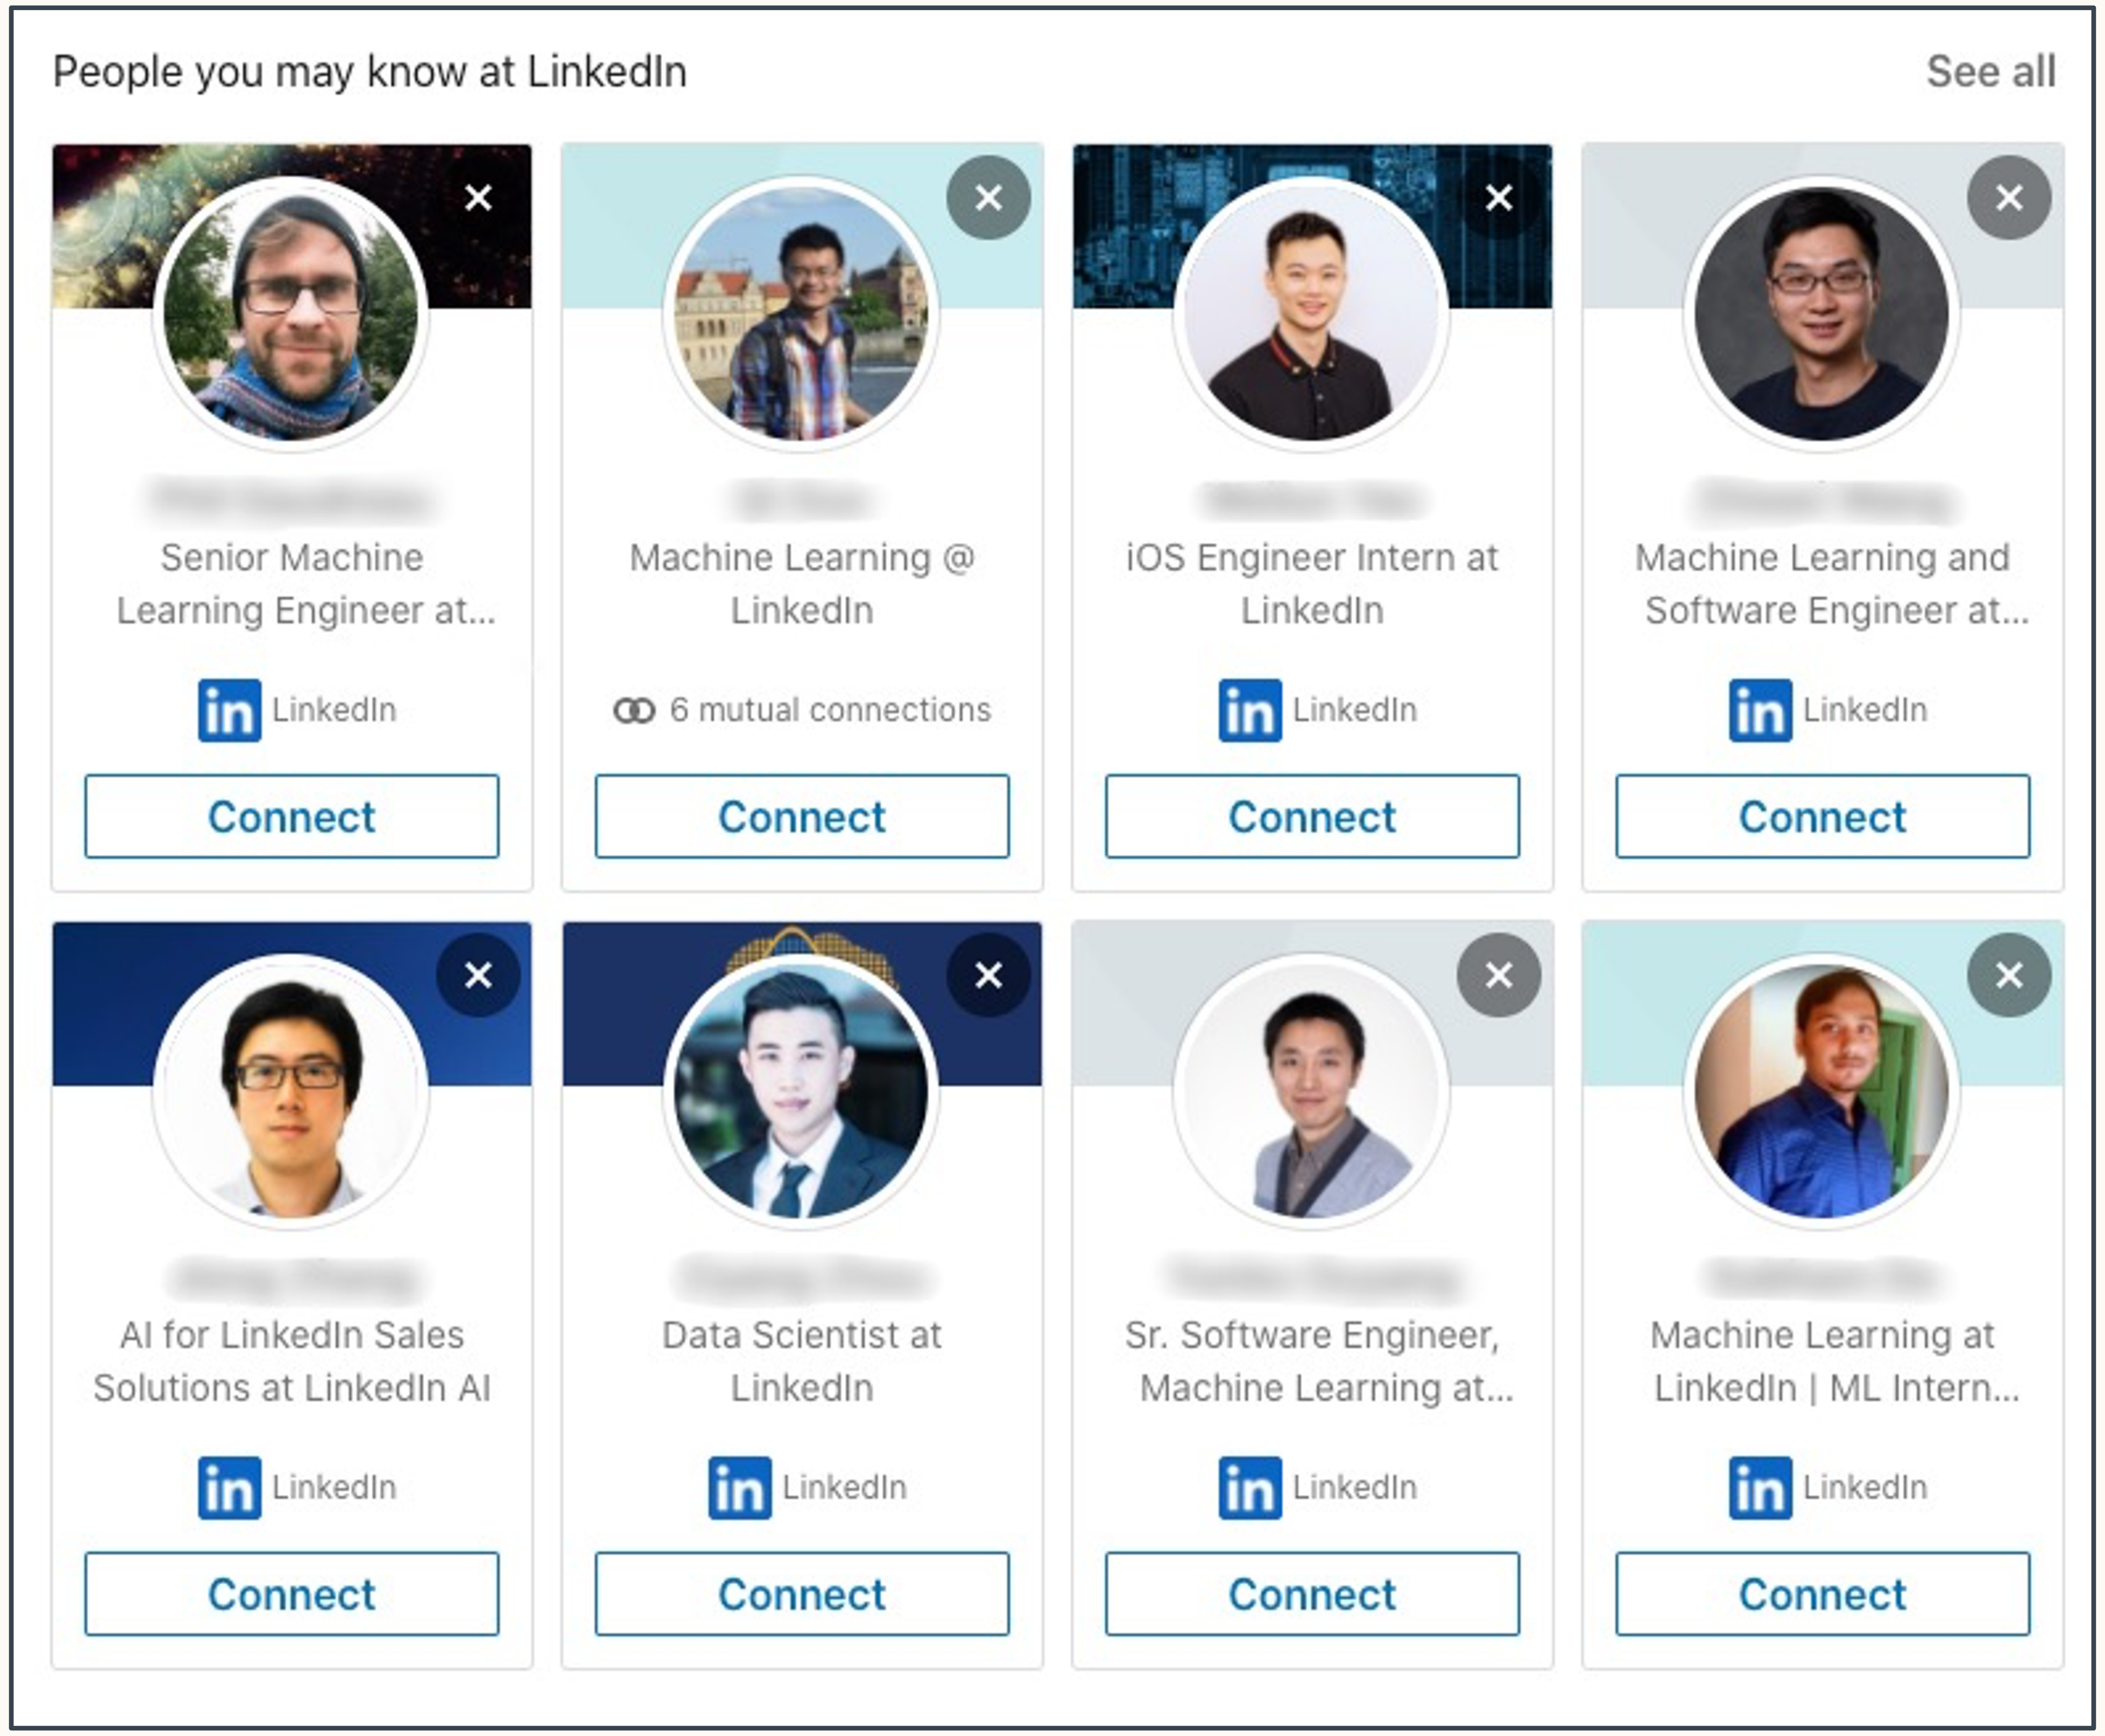

Source: Member who looks at ranking **Destination**: Members being ranked

A General Framework

Fairness of Exposure in Rankings

Ashudeep Singh
Cornell University
Ithaca, NY
ashudeep@cs.cornell.edu

Thorsten Joachims
Cornell University
Ithaca, NY
tj@cs.cornell.edu

$$
\begin{aligned}
\mathbf{P} = \operatorname{argmax}_{\mathbf{P}} \quad & \mathbf{u}^T \mathbf{P} \mathbf{v} && \text{(expected utility)} \\
\text{s.t.} \quad & \mathbb{1}^T \mathbf{P} = \mathbb{1}^T && \text{(sum of probabilities for each position)} \\
& \mathbf{P}\mathbb{1} = \mathbb{1} && \text{(sum of probabilities for each document)} \\
& 0 \leq \mathbf{P}_{i,j} \leq 1 && \text{(valid probability)} \\
& \mathbf{P} \text{ is fair} && \text{(fairness constraints)}
\end{aligned}
$$

Preliminaries

P forms a doubly stochastic matrix of size N × N

Probabilistic Rankings

$P_{i,j}$: the probability that ranking R places item **i** at rank **j**

$$P = \begin{bmatrix} 0.4 & 0.5 & 0.1 \\ 0.3 & 0.3 & 0.4 \\ 0.3 & 0.2 & 0.5 \end{bmatrix}$$

Birkhoff-von Neumann (BvN) decomposition:

doubly stochastic matrix into a convex sum of permutation matrices

$$\mathbf{A} = \theta_1 \mathbf{A_1} + \theta_2 \mathbf{A_2} + \cdots + \theta_n \mathbf{A_n}$$

$$R' = \begin{vmatrix} 0 & 0.4 & 0.2 & 0.4 \\ 0.4 & 0.4 & 0 & 0.2 \\ 0.4 & 0.2 & 0.4 & 0 \\ 0.2 & 0 & 0.4 & 0.4 \end{vmatrix} = 0.4 \begin{vmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} + 0.4 \begin{vmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{vmatrix} + 0.2 \begin{vmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{vmatrix}$$

Preliminaries

Probabilistic Rankings

P is a doubly stochastic matrix of size N × N

$$\mathbf{v}_j = \frac{1}{\log(1+j)}$$

Expected utility for a probabilistic ranking:

$$\mathrm{U}(\mathbf{P}|q) = \sum_{d_i \in \mathcal{D}} \sum_{j=1}^{N} \mathbf{P}_{i,j}\, u(d_i|q)\, v(j)$$

Rewrite it as a matrix product:

$$\mathrm{U}(\mathbf{P}|q) = \mathbf{u}^T \mathbf{P} \mathbf{v}$$

Discounted Cumulative Gain (DCG) can be represented in this format:

$$DCG(r|q) = \sum_{u \in \mathcal{U}} P(u|q) \sum_{d \in \mathcal{D}} \frac{2^{\mathrm{rel}(d|u,q)} - 1}{log(1 + \mathrm{rank}(d|r))}$$

Optimizing Fair Rankings via Linear Programming

- maximizing source side utility
- Solve problem per session!
- fairness constraints on the destination members.

P(i, j): probability of showing the i-th destination member in the j-th slot

$$
\begin{aligned}
\mathbf{P} = \operatorname{argmax}_{\mathbf{P}} \quad & \mathbf{u}^T \mathbf{P} \mathbf{v} && \text{(expected utility)} \\
\text{s.t.} \quad & \mathbb{1}^T \mathbf{P} = \mathbb{1}^T && \text{(sum of probabilities for each position)} \\
& \mathbf{P}\mathbb{1} = \mathbb{1} && \text{(sum of probabilities for each document)} \\
& 0 \le \mathbf{P}_{i,j} \le 1 && \text{(valid probability)} \\
& \mathbf{P} \text{ is fair} && \text{(fairness constraints)}
\end{aligned}
$$

Singh, Ashudeep, and Thorsten Joachims. "Fairness of exposure in rankings." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2018.

Fairness Measures

- Destination Side Metrics
  - Demographic Parity (Equality of Exposure)
    - Average exposure of the groups are equal in each query.

$$\boxed{\text{Exposure}(G_0|\mathbf{P}) = \text{Exposure}(G_1|\mathbf{P})}$$

  - Disparate Treatment
    - Average exposure of the groups proportional to their utility are equal in each query.

$$\boxed{\frac{\text{Exposure}(G_0|\mathbf{P})}{\text{U}(G_0|q)} = \frac{\text{Exposure}(G_1|\mathbf{P})}{\text{U}(G_1|q)}}$$

  - Disparate Impact
    - Average impact (eg CTR) of the groups proportional to their utility are equal in query

$$\boxed{\frac{\text{CTR}(G_0|\mathbf{P})}{\text{U}(G_0|q)} = \frac{\text{CTR}(G_1|\mathbf{P})}{\text{U}(G_1|q)}}$$

Demographic Parity

Position Bias

$$\mathbf{v}_j = \frac{1}{\log(1+j)}$$

Avg Exposure for a group

$$\text{Exposure}(G_k|\mathbf{P}) = \frac{1}{|G_k|}\sum_{d_i \in G_k} \text{Exposure}(d_i|\mathbf{P})$$

Exposure for a member

$$\text{Exposure}(d_i|\mathbf{P}) = \sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j$$

Average exposure of the groups are equal in each query.

$$\text{Exposure}(G_0|\mathbf{P}) = \text{Exposure}(G_1|\mathbf{P})$$

$$\Leftrightarrow \frac{1}{|G_0|}\sum_{d_i \in G_0}\sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j = \frac{1}{|G_1|}\sum_{d_i \in G_1}\sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j$$

$$\Leftrightarrow \sum_{d_i \in \mathcal{D}}\sum_{j=1}^{N}\left(\frac{\mathbb{1}_{d_i \in G_0}}{|G_0|} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|}\right)\mathbf{P}_{i,j}\mathbf{v}_j = 0$$

$$\Leftrightarrow \mathbf{f}^T P\mathbf{v} = 0 \qquad \left(\text{with } \mathbf{f}_i = \frac{\mathbb{1}_{d_i \in G_0}}{|G_0|} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|}\right)$$

Disparate Treatment

$$\text{Exposure}(G_k|\mathbf{P}) = \frac{1}{|G_k|} \sum_{d_i \in G_k} \text{Exposure}(d_i|\mathbf{P})$$

Average exposure of the groups proportional to their utility are equal in each query.

$$\text{U}(G_k|q) = \frac{1}{|G_k|} \sum_{d_i \in G_k} \mathbf{u}_i$$

$$\frac{\text{Exposure}(G_0|\mathbf{P})}{\text{U}(G_0|q)} = \frac{\text{Exposure}(G_1|\mathbf{P})}{\text{U}(G_1|q)}$$

$$\Leftrightarrow \frac{\frac{1}{|G_0|} \sum_{d_i \in G_0} \sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j}{\text{U}(G_0|q)} = \frac{\frac{1}{|G_1|} \sum_{d_i \in G_1} \sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j}{\text{U}(G_1|q)}$$

$$\Leftrightarrow \sum_{i=1}^{N} \sum_{j=1}^{N} \left(\frac{\mathbb{1}_{d_i \in G_0}}{|G_0|\text{U}(G_0|q)} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|\text{U}(G_1|q)} \right) \mathbf{P}_{i,j}\mathbf{v}_j = 0$$

$$\Leftrightarrow \mathbf{f}^T P \mathbf{v} = 0 \qquad (\text{with } \mathbf{f}_i = \left(\frac{\mathbb{1}_{d_i \in G_0}}{|G_0|\text{U}(G_0|q)} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|\text{U}(G_1|q)} \right))$$

Disparate Impact

$$
\begin{aligned}
P(\text{click on document } i) &= P(\text{examining } i) \times P(i \text{ is relevant}) \\
&= \text{Exposure}(d_i|\mathbf{P}) \times P(i \text{ is relevant}) \\
&= \left(\sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{v}_j\right) \times \mathbf{u}_i
\end{aligned}
$$

Average impact (e.g. CTR) of the groups proportional to their utility are equal in query

$$
\text{CTR}(G_k|\mathbf{P}) = \frac{1}{|G_k|} \sum_{i \in G_k} \sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{u}_i\mathbf{v}_j
$$

$$
\begin{aligned}
& \frac{\text{CTR}(G_0|\mathbf{P})}{\text{U}(G_0|q)} = \frac{\text{CTR}(G_1|\mathbf{P})}{\text{U}(G_1|q)} \\
\Leftrightarrow\ & \frac{\frac{1}{|G_0|}\sum_{i \in G_0}\sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{u}_i\mathbf{v}_j}{\text{U}(G_0|q)} = \frac{\frac{1}{|G_1|}\sum_{i \in G_1}\sum_{j=1}^{N} \mathbf{P}_{i,j}\mathbf{u}_i\mathbf{v}_j}{\text{U}(G_1|q)} \\
\Leftrightarrow\ & \sum_{i=1}^{N}\sum_{j=1}^{N}\left(\frac{\mathbb{1}_{d_i \in G_0}}{|G_0|\text{U}(G_0|q)} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|\text{U}(G_1|q)}\right)\mathbf{u}_i\mathbf{P}_{i,j}\mathbf{v}_j = 0 \\
\Leftrightarrow\ & \mathbf{f}^T P\mathbf{v} = 0 \qquad\qquad (\text{with } \mathbf{f}_i = \left(\frac{\mathbb{1}_{d_i \in G_0}}{|G_0|\text{U}(G_0|q)} - \frac{\mathbb{1}_{d_i \in G_1}}{|G_1|\text{U}(G_1|q)}\right)\mathbf{u}_i)
\end{aligned}
$$

Two-sided Fair Marketplace

optimization-based framework

- maximizing source side utility
- maintaining fairness constraints on the destination members.
- post-processing
- Detailed Design

Expanded framework

- Fairness constraints for destination members at multiple sessions by considering the multi-session utility
- Adding fairness constraints for source members
- Post-processing
- Detailed Design

$$\text{Maximize } u_s^\top P_s v$$

$$\text{Subject to } 1^\top P_s = 1, \quad P_s 1 \leq 1, \quad 0 \leq P_s(d,r) \leq 1, \quad \underbrace{\boxed{f_{k,k'}^\top P_s w = 0}}_{\text{Destination Side Fairness}} \text{ and } \underbrace{\boxed{\tilde{u}_s P_s w = c}}_{\text{Multi-Session Fairness}}$$

Two-Sided Fair Marketplace

A Post-Processesing approach:

- Operates on the final ranking with fairness constraints (very generally applicable)
- Agnostic to the underlying model architecture
- Suitable for large-scale recommender systems

Two-sided Fair Marketplace

Most recommender systems:

maximizing source side utility

$$U_q^{\text{source}}(s) = \sum_{d=1}^{D_q} \sum_{r=1}^{m} u_{s,d}^q P_s^q(d,r) v_r = u_s^\top P_s v$$

P(d, r): probability of showing the d-th destination member in the r-th slot for source s and query q

Expanded Framework:

Destination side utility

1. much harder to tackle
2. depends of different queries arising from different sources

$$U_q^{\text{dest}}(s,d) = u_{s,d}^q \sum_{r} P_s^q(d,r) \cdot v_r = u_{s,d}^q P_s^q(d,\cdot) v$$

Destination Side Utility

expected utility a *destination* member d receives

$$U^{\text{dest}}(d)[T] = \sum_{i=1}^{N} \rho^{T-t_i} U^{\text{dest}}_{q_{t_i}}(s_{t_i}, d)$$

ρ is an appropriately chosen discount rate

average cumulative utility for group G at time t

$$\mu_{G_k}[t] = \frac{1}{\underbrace{|G_k|}_{\text{gender group}}} \sum_{d \in G_k} U^{dest}(d)[t]$$

incremental utility

$$\Delta\mu_{G_k}[T] = \mu_{G_k}[T] - \mu_{G_k}[t]$$

equality of average incremental utility across groups
(Multi-session Dest-side Fairness)

$$\Delta\mu_{G_k}[T] = \Delta\mu_{G_{k'}}[T]$$

Multi-session Destination Side Fairness Constraint

$$\boxed{\Delta\mu_{G_k}[T] = \Delta\mu_{G_{k'}}[T]}$$

We can rewrite it as:

$$\boxed{\tilde{u}_s P_s w = c}$$

$$\boxed{u^q_{s,d}\left(\frac{\mathbb{1}_{d\in G_k}}{|G_k|} - \frac{\mathbb{1}_{d\in G_{k'}}}{|G_{k'}|}\right)}$$

$$\boxed{c = (1-\rho^{T-t})(\mu_{G_k}[t] - \mu_{G_{k'}}[t])}$$

Two-sided Fair Marketplace Optimization Details

Optimization problem with multi-session Fairness:

Maximize $u_s^\top P_s v$

Subject to $1^\top P_s = 1, \quad P_s 1 \leq 1, \quad 0 \leq P_s(d,r) \leq 1, \quad \underbrace{\boxed{f_{k,k'}^\top P_s w = 0}}_{\text{Destination Side Fairness}}$ and $\underbrace{\boxed{\tilde{u}_s P_s w = c}}_{\text{Multi-Session Fairness}}$

Vectorized form:

Max $p^\top (u \cdot v)$

s.t. $\quad p^\top (f_{k,k'} \cdot v) = 0, \;\; p^\top (\tilde{u} \cdot v) = c, \;\; [I_m : I_m : \cdots : I_m] p = 1$ and $p_d \in T_m \;\; \forall\, d \in M$

Add a regularizer + better format

Maximize $p^\top (u \cdot v) + \frac{\gamma}{2} p^\top p \qquad$ subject to $\quad Ap \leq b \quad$ and $\quad p_d \in T_m \;\; \forall\, d \in M$

Two-sided Fair Marketplace Optimization Details

$$\textbf{Maximize } p^\top (u \cdot v) + \frac{\gamma}{2} p^\top p \qquad \textbf{subject to} \quad Ap \leq b \quad \textbf{and} \quad p_d \in T_m \quad \forall\, d \in M$$

- Solve the optimization problem
- Obtain the dual variables (λ_1, λ_2, η)
- Obtain the primal solution from duals
  (Drastically reduces the latency cost)

$$\hat{p}_d(\lambda_1, \lambda_2, \eta) = \Pi_{T_m}\left(\frac{1}{\gamma}\{(u \cdot v)_d - \lambda_1(f_{k,k'} \cdot v)_d - \lambda_2(\tilde{u} \cdot v)_d - \eta\}\right)$$

Extension to Source Side Fairness

ensuring that they receive comparable expected utility across groups over multiple sessions.

$$\mathbb{E}\left[U^{\text{source}}(G_k, T)\right] = \sum_{i:t_i \leq T} \rho^{T-t_i} U^{\text{source}}_{q_i}(s_i) = \sum_{i:t_i \leq T} \rho^{T-t_i} u_{s_i}^{\top} P_{s_i} v$$

$$\mathbb{E}\left[U^{\text{source}}(G_k, T+\delta)\right] = \mathbb{E}\left[U^{\text{source}}(G_{k'}, T+\delta)\right]$$

$$u_s^{\top} P_s v = \mathbb{E}\left[U^{\text{source}}(G_{k'}, T+\delta)\right] - \rho^{\delta}\mathbb{E}\left[U^{\text{source}}(G_k, T)\right]$$

since the ranking algorithms are by default trying to optimize the source side utility, it is unlikely that a group of source members would be consistently achieving poorer utility in a well- functioning system

Extension to Source Side Fairness

ensuring that they receive comparable expected utility across groups over multiple sessions.

$$\text{Minimize} \sum_{k' \neq k} \left| u_s^\top P_s v - c_{k,k'} \right|$$

$$\text{Subject to } 1^\top P_s = 1, \quad P_s 1 \leq 1, \quad 0 \leq P_s(d,r) \leq 1, \quad f_{k,k'}^\top P_s w = 0, \quad \tilde{u}_s P_s w = c$$

$$c_{k,k'} = \mathbb{E}\left[U^{\text{source}}(G_{k'}, T+\delta)\right] - \rho^\delta \mathbb{E}\left[U^{\text{source}}(G_k, T)\right]$$

since the ranking algorithms are by default trying to optimize the source side utility, it is unlikely that a group of source members would be consistently achieving poorer utility in a well- functioning system

Experiments (Simulation Setup)

- A graph to represent connections between members
- Standard undirected graph, not bipartite
- A_t is the adjacency matrix at time t

$$\boxed{\mathrm{Aff}_{\mathrm{Graph}}(i,j) \;=\; A_t^2(i,j)}$$

$$\boxed{\mathrm{Aff}_{\mathrm{member}}(i,j) \;=\; -\|X_i - X_j\|_2}$$

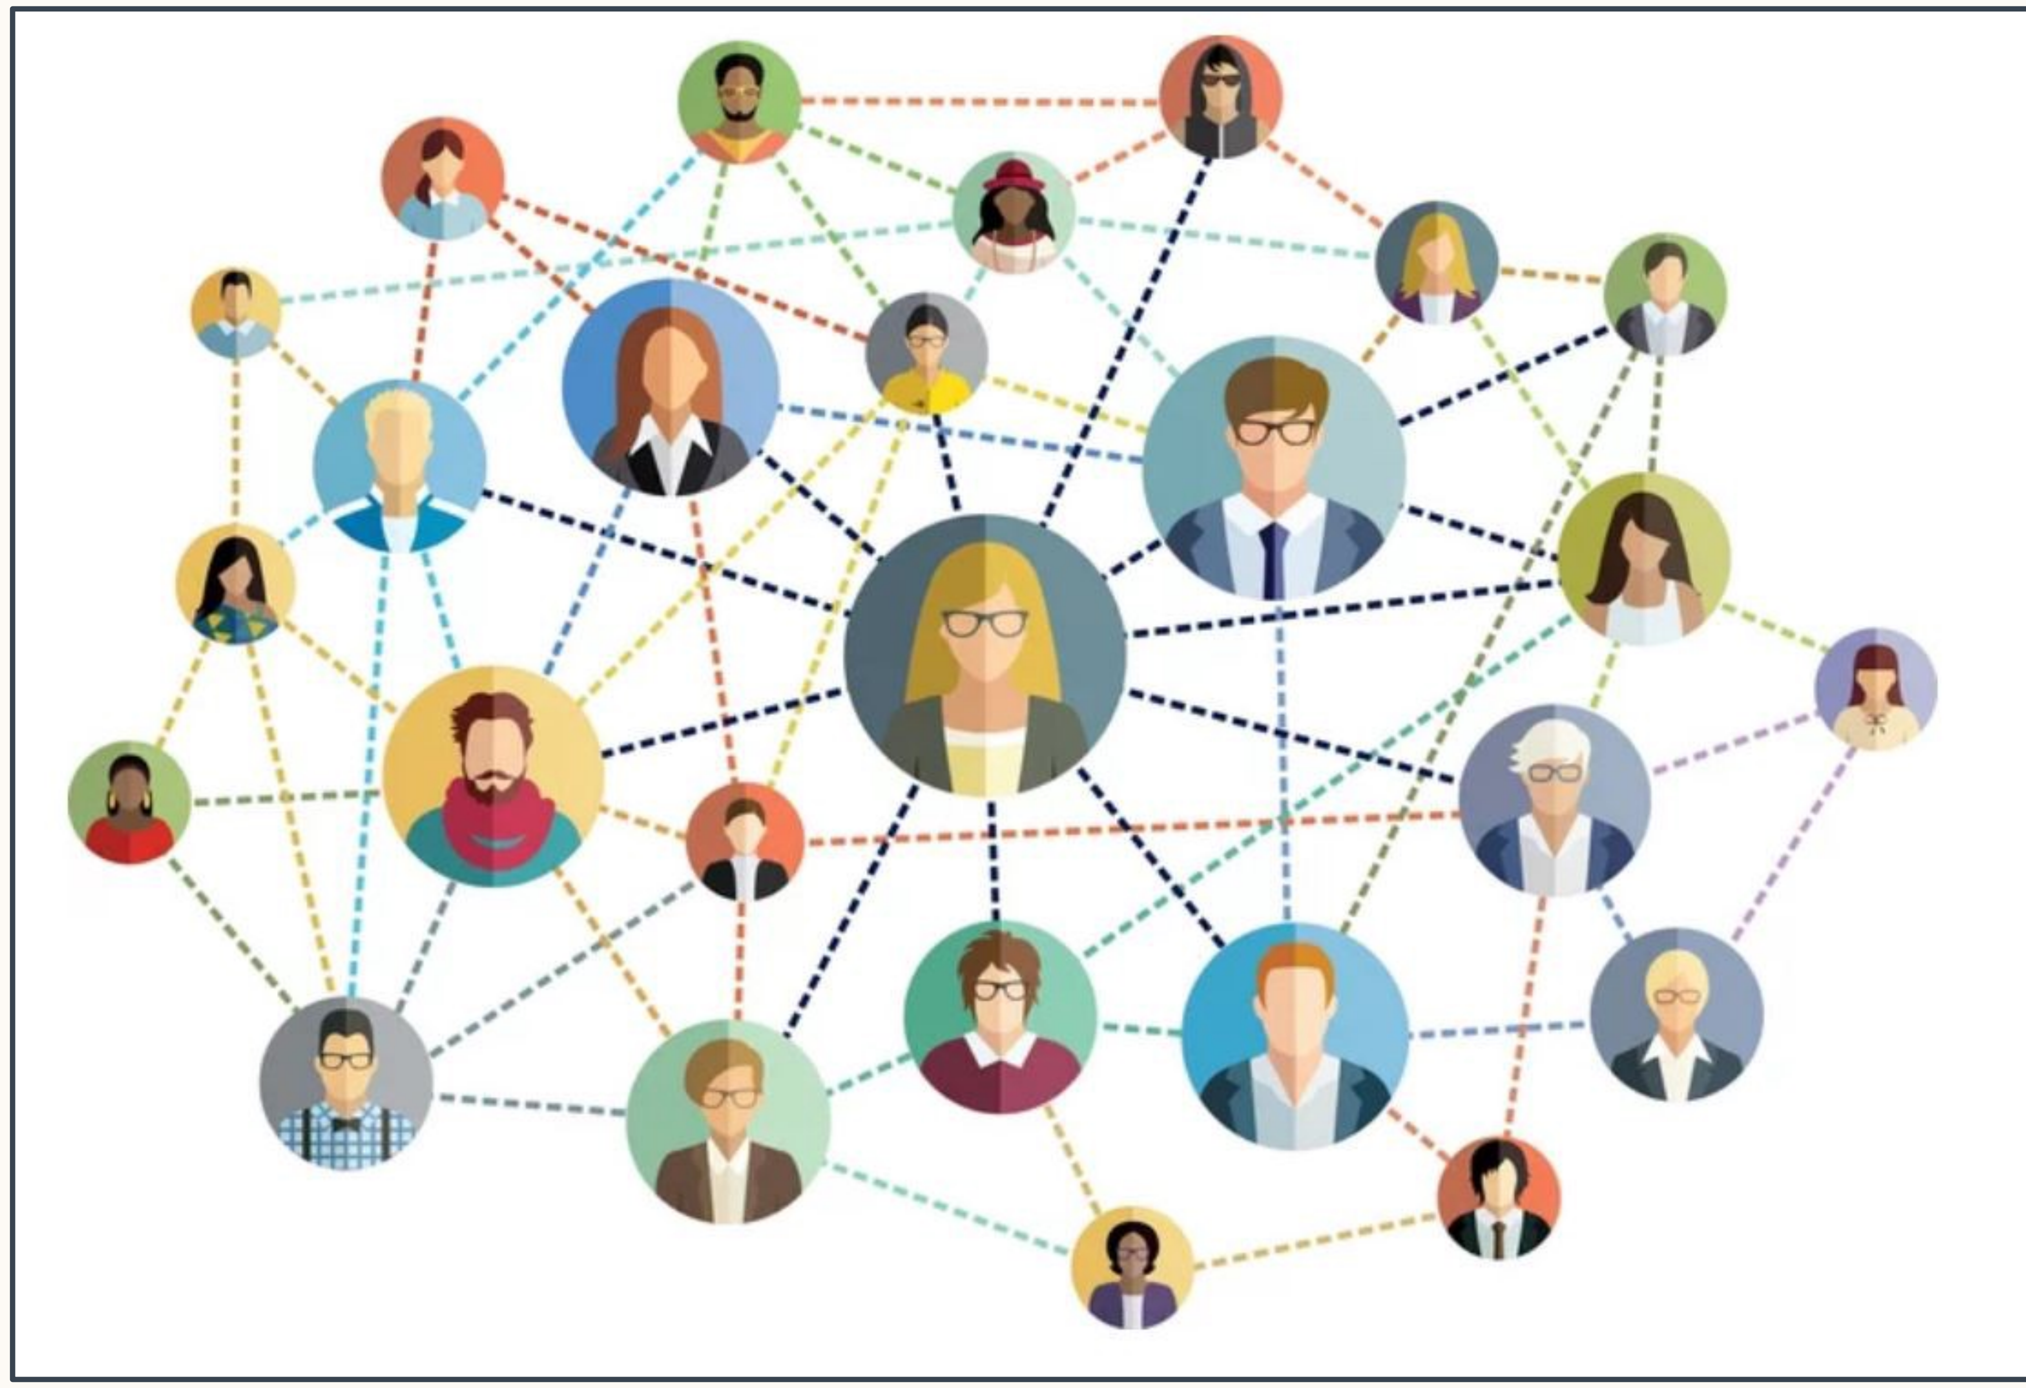

- Graph Evolution

A vertex i is picked at random.
For each other vertex j, a "model" score s(i, j) is computed as

$$\boxed{s(i,j) = p(G_i, G_j)\exp(\lambda \mathrm{Aff}_{\mathrm{Graph}}(i,j))\exp(\mu \mathrm{Aff}_{\mathrm{member}}(i,j))}$$

Experiments

Fairness adjustments:

1. No fairness adjustment in the ranking (noReranker)
2. Equality of opportunity, using the primal for destination side fairness (primal)
3. Equality of opportunity, using the dual method for destination side fairness (dualNoDynamic)
4. Multi-Session destination utility adjustment via the dual (dualWithDynamic).

Simulation Details:
1000 members and 1000 iterations

Experiments (Simulation Results)

Settings	Method	ΔDP	$\Delta\lvert\text{DP}\rvert$	Ratio of Source Utility	% Total Destination Utility

shown members
$m = 10$
$t = 50$
rate of dual refresh
	`noReranker`
	`primal`
	`dualNoDynamic`
	`dualWithDynamic`
		-0.0216
		0.0179
		0.0055
		0.0249
			0.1964
			0.1671
			0.2125
			0.2008
				0.5100
				0.5024
				0.5100
				0.5100
					0.7051
					0.5371
					0.5860
					0.5706

$m = 20$
$t = 50$
	`noReranker`
	`primal`
	`dualNoDynamic`
	`dualWithDynamic`
		-0.0072
		0.0109
		0.0122
		0.0155
			0.1182
			0.1002
			0.1090
			0.1099
				0.5100
				0.5023
				0.5099
				0.5099
					0.7059
					0.5269
					0.5536
					0.5371

$m = 20$
$t = 20$
	`noReranker`
	`primal`
	`dualNoDynamic`
	`dualWithDynamic`
		-0.0072
		0.0109
		0.0177
		0.0185
			0.1182
			0.1002
			0.1105
			0.1090
				0.5100
				0.5023
				0.5099
				0.5099
					0.7059
					0.5269
					0.5210
					0.5166

Adding fairness re-rankers:

- help rebalance *destination* utility (last column)
- not significantly altering *source* utility metrics (see the 5th column)

Conclusion (My Thoughts)

1. Novel Idea by Designing multi-session dest-side utility
   Enables introducing multi-session dest-side utility and source-side utility

2. Using dual trick makes it flexible for large-scale implementation

3. Evaluation on only one fairness constraint (Demographic Parity)
   Since utility is available in test time Disparate Treatment could be evaluated too.

4. Does not provide different trade-offs between Utility and Fairness

5. Evaluation needed on Real-world Large-Scale data

Thank you