Personalizing Fairness-aware Re-ranking

Weiwen Liu[1], Robin Burke[2]

[1]The Chinese University of Hong Kong
[2]Depaul University

Multi-sided Recommender Systems (MRS)

Users/Consumers are not the only stakeholder in some recommendation scenarios.

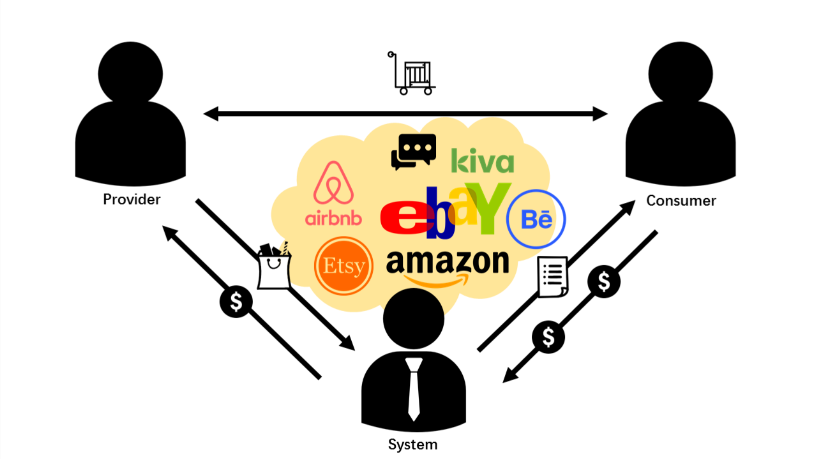

Multi-sided Recommender Systems (MRS)

Consumer: Expect personalized recommendations to meet their interests and needs.

Provider: Offer items to the system and benefit from consumer choices.

System: Receive items from providers and recommend them to the consumers.

Multi-sided Recommender Systems (MRS)

Consumer: Expect personalized recommendations to meet their interests and needs.

Provider: Offer items to the system and benefit from consumer choices.

System: Receive items from providers and recommend them to the consumers.

Provider-side Fairness

Twelve animals wait to become superstars

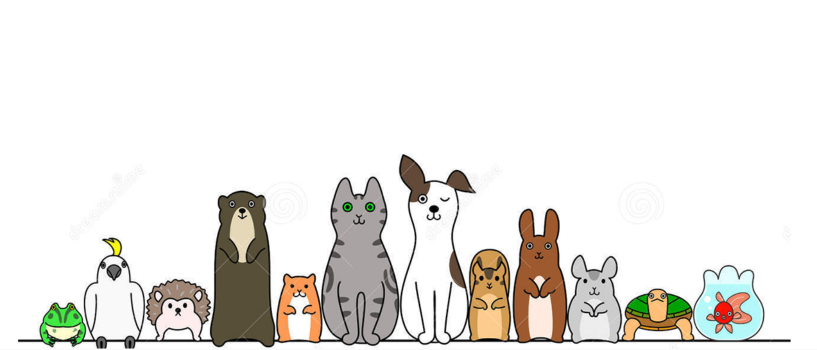

Provider-side Fairness

Only two of them get exposed...

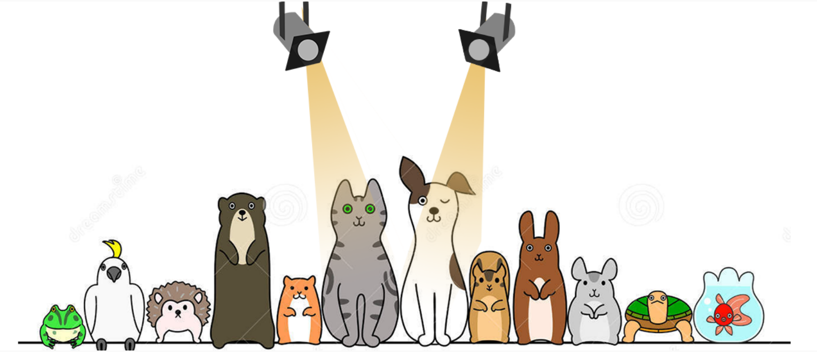

Provider-side Fairness: Kiva.org

Kiva.org is a non-profit site for crowd-sourced micro-lending.

Provider-side Fairness: Kiva.org

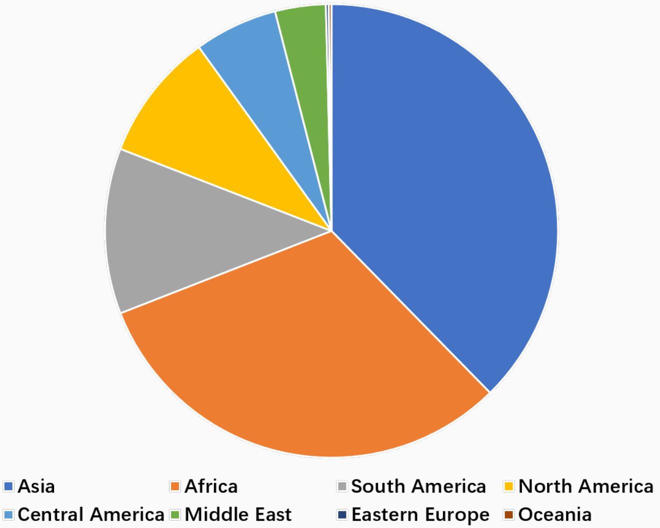

Figure 1: Number of recommendations for each region. WRMF.

Provider-side Fairness

Provider

- passiveness
- competitiveness
- a key role in MRS

Goal

To balance across different providers rather than concentrating on certain dominant ones

Problem Formulation

- Given a set of users $U = \{1, \ldots, m\}$, a set of items $V = \{1, \ldots, n\}$ and an initial ranking list $R = [1, \ldots, z]$.

Problem Formulation

- Given a set of users $U = \{1, \ldots, m\}$, a set of items $V = \{1, \ldots, n\}$ and an initial ranking list $R = [1, \ldots, z]$.
- Each provider $d \in D$ owns a set of items to be recommended.

Problem Formulation

- Given a set of users $U = \{1, \ldots, m\}$, a set of items $V = \{1, \ldots, n\}$ and an initial ranking list $R = [1, \ldots, z]$.
- Each provider $d \in D$ owns a set of items to be recommended.
- Our task is to generate a re-ranked list S of K distinct items that is both accurate and fair.

Algorithm

We designed a re-ranking criterion:

$$\overbrace{P(v|u)}^{\text{accuracy}} + \overbrace{\lambda \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}}}^{\text{fairness}},$$

- When recommending item list to user u, item that scores most is added to the output list S repeatedly $\Rightarrow$ a greedy algorithm

Algorithm

We designed a re-ranking criterion:

$$\overbrace{P(v|u)}^{\text{accuracy}} + \overbrace{\lambda \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}}}^{\text{fairness}},$$

- When recommending item list to user u, item that scores most is added to the output list S repeatedly $\Rightarrow$ a greedy algorithm
- For an item $v \in d$, if S does not cover d, then an additional positive term will be added to the estimated user preference $P(v|u)$ $\Rightarrow$ to balance between accuracy and fairness

Algorithm

We designed a re-ranking criterion:

$$\overbrace{P(v|u)}^{\text{accuracy}} + \overbrace{\lambda \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}}}^{\text{fairness}},$$

- When recommending item list to user u, item that scores most is added to the output list S repeatedly $\Rightarrow$ a greedy algorithm
- For an item $v \in d$, if S does not cover d, then an additional positive term will be added to the estimated user preference $P(v|u)$ $\Rightarrow$ to balance between accuracy and fairness
- As formulated, the criterion favors the items that belong to multiple providers.

Algorithm

Algorithm 1 Fairness-Aware Re-ranking

Input: u, R, K, λ, τ_u

Output: S

$S \leftarrow \emptyset$
while $|S| < K$ **do**
    $v^* \leftarrow \arg\max_{v \in R \setminus S} P(v|u) + \lambda \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}}$
    $R \leftarrow R \setminus \{v^*\}$
    $S \leftarrow S \cup \{v^*\}$
end while
return S

Diversity Tolerance

The tolerance towards exploration or diversification in providers varies for different consumers.

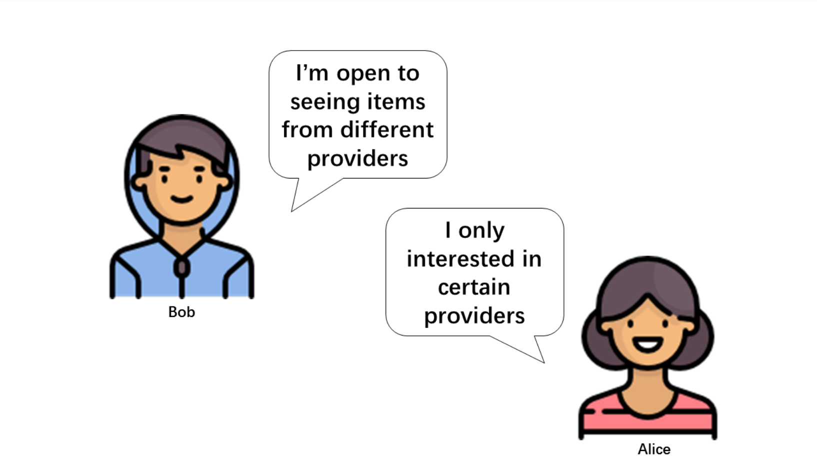

Diversity Tolerance

The user tolerance towards different providers τ_u is defined by

$$\tau_u = -\sum_{d \in D} I(d|u) \log I(d|u),$$

$$I(d|u) = \frac{\sum_v r(u,v) \mathbb{1}_{\{v \in d\}}}{\sum_v \sum_{d' \in D} r(u,v) \mathbb{1}_{\{v \in d'\}}},$$

$r(u,v)$ is the rating from user u to item v.

Fairness-Aware Re-ranking (FAR)

$$P(v|u) + \lambda \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}},$$

Personalized Fairness-Aware Re-ranking (PFAR)

$$P(v|u) + \lambda \tau_u \sum_{d \in D} P(d) \mathbb{1}_{\{v \in d\}} \prod_{i \in S} \mathbb{1}_{\{i \notin d\}}, \tag{1}$$

Evaluation Criterion

- Average Provider Coverage Rate (APCR)

$$\text{APCR} = \frac{1}{|U_t|} \sum_{u \in U_t} \frac{\#recommended\_provider}{\#provider}, \tag{2}$$

  where U_t is the test user set.

- APCR@NDCG$_{5\%}$
  The increased APCR value obtained when we allow a 5% decrease in nDCG.

Experiment: Synthetic Data

Table 1: Movielens: APCR@NDCG$_{5\%}$ on the data sets (%), providers assigned at random.

	PFAR		FAR	
	λ	APCR@NDCG$_{5\%}$	λ	APCR@NDCG$_{5\%}$
itemKNN	2.45	77.47 (+70.24)	1.62	78.33 (+72.14)
userKNN	0.15	67.89 (+57.94)	0.10	68.98 (+60.48)
rankALS	0.29	72.81 (+59.01)	0.22	74.43 (+62.57)
WRMF	0.24	70.45 (+58.79)	0.16	72.10 (+62.50)

Experiment: Kiva.org

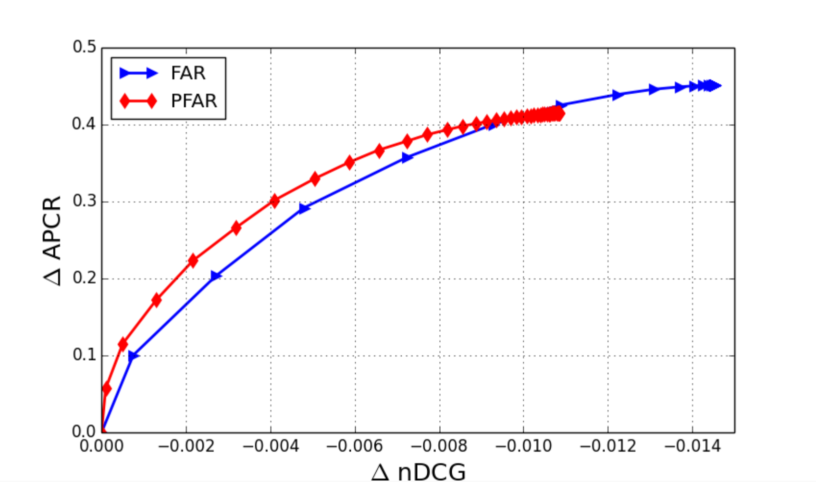

Figure 2: Change in nDCG and APCR with increasing λ (range 0 to 2.0 in steps of 0.05).

Experiment: Kiva.org

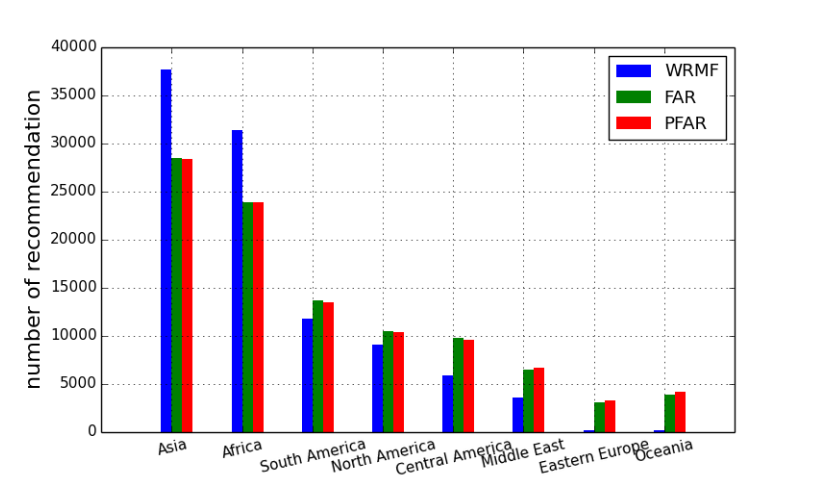

Figure 3: Number of recommendations for each region.

Conclusion

- We formulate a recommendation scenario in a multi-sided recommender system and define the fairness requirement for providers.
- We design a re-ranking algorithm to balance between personalization and fairness, and propose the incorporation of diversity tolerance of individuals.
- We show the results of experiments conducted on synthetic and real-world data to validate the effectiveness of our proposed algorithm.

Future Work

- Explore different methods for computing personalized diversity tolerance factors, e.g to solve the cold start problem
- Examine variants of the re-ranking algorithm to take into account the size of each providers' inventory.
- Adjust the accuracy/coverage tradeoff in a dynamic way as items are ranked, valuing accuracy more at the top of the list and provider coverage more at the bottom of the list.
- Explore the online fairness algorithm.