

# Multi-stakeholder Recommendation and its Connection to Multi-sided Fairness\*

Himan Abdollahpouri  
University of Colorado Boulder, USA  
himan.abdollahpouri@colorado.edu

Robin Burke  
University of Colorado Boulder, USA  
robin.burke@colorado.edu

## ABSTRACT

There is growing research interest in recommendation as a multi-stakeholder problem, one where the interests of multiple parties should be taken into account. This category subsumes some existing well-established areas of recommendation research including reciprocal and group recommendation, but a detailed taxonomy of different classes of multi-stakeholder recommender systems is still lacking. Fairness-aware recommendation has also grown as a research area, but its close connection with multi-stakeholder recommendation is not always recognized. In this paper, we define the most commonly observed classes of multi-stakeholder recommender systems and discuss how different fairness concerns may come into play in such systems.

## KEYWORDS

Recommender systems; Fairness in recommendation; Multi-stakeholder recommendation; Multi-sided fairness;

## 1 INTRODUCTION

Recommender systems (RS) have been used in a variety of different domains to help users find relevant and interesting items. They recommend movies to watch [14], songs to listen [9], jobs to take [17] or even a person to date [19].

The research in RS has been mainly focused on the personalization. That is, delivering the recommendations that best match the needs and interests of the end user. That is indeed an important consideration as users are one of the most important stakeholders in any recommendation platform but not the only one [1]. There are numerous examples of recommender systems in which there are other stakeholders that their needs and preferences should be taken into account. The incorporation of the objectives and preferences of different stakeholders in the recommendation process is referred to as *multi-stakeholder recommendation* [7].

Multi-stakeholder recommendation is a relatively new topic (at least in academia) and, therefore, there is still no clear understanding of what type of recommender system is a multi-stakeholder one. Defining different classes of multi-stakeholder recommendation would help to distinguish among different recommendation problems, which is important for developing the right algorithmic solutions for these types of recommender systems. Moreover, fairness-aware recommendation as a growing sub-topic in recommendation research has gained a lot of attention in recent years. Fairness in recommendation can be defined in different ways, as

\*Copyright 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).  
Presented at the RMSE workshop held in conjunction with the 13th ACM Conference on Recommender Systems (RecSys), 2019, in Copenhagen, Denmark.  
RMSE workshop at ACM RecSys 2019, September 20, Copenhagen, Denmark.

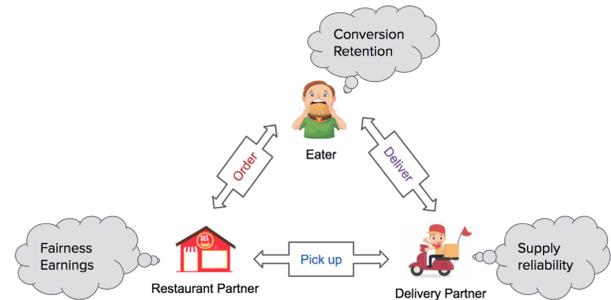


Figure 1: A multi-stakeholder problem at Uber Eats <sup>2</sup>

discussed in [24]. One of its accepted definitions is to avoid giving discriminatory recommendations to different users based on some sensitive features such as gender or race [12]. Talking about the fairness of a recommendation can be meaningful only when there are different stakeholders involved. Thus, there is a close connection between the requirements of multi-stakeholder recommendation and those of fairness-aware recommendation.

In this paper, we present three most commonly observed classes of multi-stakeholder recommendation:

- Multi-receiver recommendation
- Multi-provider recommendation
- Recommendation with side stakeholders

In addition, we discuss the fairness concerns that could come into play in these types of recommendation systems with regard to different stakeholders such as users, item providers, etc.

## 2 MULTI-STAKEHOLDER RECOMMENDATION

Recommender systems are often multi-stakeholder environments [2, 7]: there are several stakeholders (often with conflicting preferences) whose needs and preferences should be taken into account in generating the recommendations. For example, in a food delivery system like the one for Uber Eats, we can observe three major stakeholders as shown in Figure 1: 1) the eaters or otherwise users who use the app and receive the recommendations for different restaurants, 2) the restaurant that are being recommended and 3) the delivery partners that take the food from the restaurants to the user's address. Eaters want to get relevant and interesting recommendations; restaurants want to be given a fair exposure to different users so they can have enough customers and finally the delivery partners need to be satisfied with the type of orders they need to deliver. Uber Eats cannot survive without the existence of

<sup>2</sup><https://eng.uber.com/uber-eats-recommending-marketplace/>

any of these three parties and therefore it needs to take all of these stakeholders' preferences into account. This example is clearly beyond the standard user-focused recommender system in which the needs and preferences of the end user is the only consideration.

Although the research in recommender system has been mainly focused on the satisfaction of the end user, there exist several threads of research in the recommender system that are examples of multi-stakeholder recommendation but similarities among them have not typically been recognized. For example, in a group recommendation platform [10] the system wants to recommend an item or a list of items to a group of users such that it meets the preferences of the group members. That is different from recommending items to only one user as, in many cases, the preferences of the users in the group could conflict with each other and, therefore, it is important for the recommender system to find ways to handle such conflicts.

Another area of work that also falls into the multi-stakeholder definition is the reciprocal recommendation where the preferences of both the receiver and the provider of the recommendations should be taken into account [25]. For example in online dating applications [20, 22], recommending a user to another user is only meaningful if both users are happy with this recommendation.

Multi-stakeholder recommendation can be seen as an umbrella for many different types of recommendation problems involving two or more different stakeholders. These types of systems have not yet been studied in a systematic way and there is still a lot of room for further research and development.

In next section, we define the most commonly observed classes of multi-stakeholder recommendation and provide examples for each. We believe having a well-defined architecture for these systems can facilitate further research in this area.

### 3 CLASSES OF MULTI-STAKEHOLDER RECOMMENDATION

In this section we define different classes of multi-stakeholder recommendation according to the architecture of the recommendation platform. One could argue there exist more types of architectures but we believe these are typically the main and commonly observed patterns.

#### 3.1 Multi-receiver Recommendation

The first class of multi-stakeholder recommendation is *Multi-receiver Recommendation* where the receiver of the recommendation is not one individual but rather a group of individuals. In these systems, the recommendation should be appealing to all the individuals receiving the recommendation. Depending on the characteristics of the receivers of the recommendation, there could exist two types of multi-receiver recommendation:

- **Heterogeneous:** In this type of multi-receiver recommendation, the receivers of the recommendation may not necessarily be the same type. For example, in an educational system that recommends courses to students, the receivers of the recommendations are the students who take those courses. In addition, in many cases, the parents of those students could be also considered as the receiver of the recommendation even though they are not taking the course themselves but

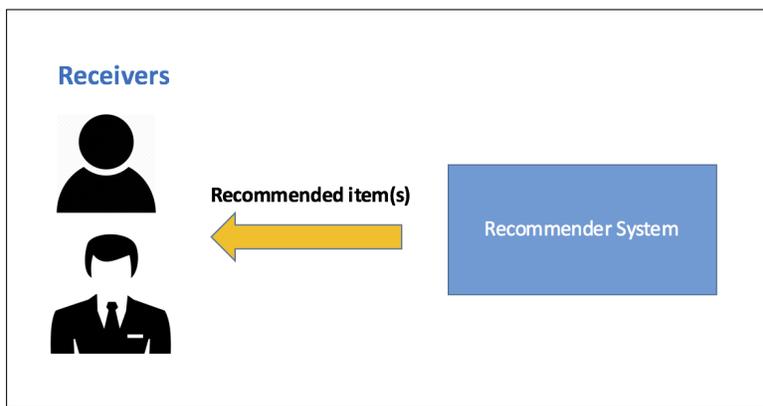
they somehow involve in the challenges and requirements that comes with those recommended course (e.g. paying the fee for the registration etc.) so their preferences should be also taken into account in recommending those courses. Zheng et al. [27] proposed a utility-based multi-stakeholder recommendations to the area of personalized learning in educations in which the preferences of students and instructors are taken into account simultaneously. Figure 2-a shows a typical architecture for this type of systems. We showed the users with different icons to clarify the heterogeneity of the receivers.

- **Homogeneous:** In contrast to the heterogeneous multi-receiver recommendation, users in homogeneous multi-receiver recommendation are all from the same type (e.g. students) and they are consuming the recommended item. This type of multi-receiver recommendation is known as *group recommendation* in the literature where an item or set of items is recommended to a group of users who all use or consume the recommendation. For example, a movie recommender system that recommends movies to a group of friends or family members falls within this category as all these users are the real consumers of these recommendations. See the survey in [15]. Figure 2-b shows a typical architecture for this type of systems. The users are all shown with the same icon.

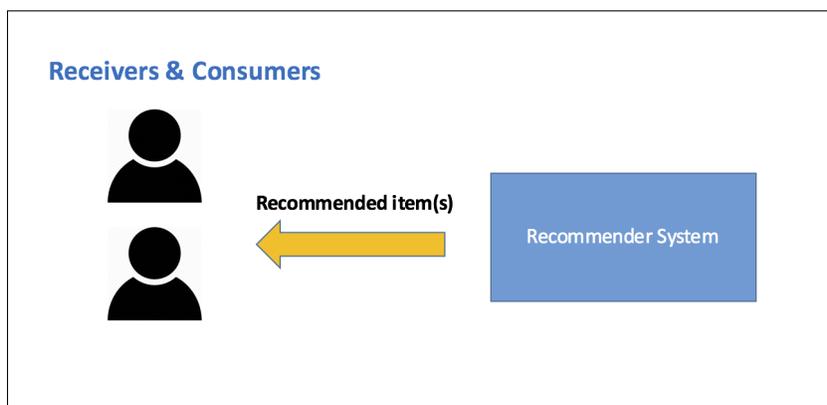
#### 3.2 Multi-provider Recommendation

It is often the case that the items and products on a recommender system are provided by several parties that use the platform for reaching out to their desired audience. For example, in a sharing economy platform like Airbnb all the available listings on the platform are actually provided by different hosts and the Airbnb itself does not own any of these apartments. These hosts are using the Airbnb platform to find travellers that could be interested in staying at their room or apartment. Therefore, Airbnb is a *multi-provider recommendation* platform since there are numerous number of hosts in the system and Airbnb should try to give a fair amount of exposure to the listings provided by each of these hosts. Several interesting challenges could arise from this type of systems such as cold start providers (a provider that is new to the system and needs to get attention), malicious providers (a provider who is violating some rules or s/he has bad ratings from the users) and problems related to fairness which we discuss later in section 4. Figure 3 shows the architecture for this type of systems. As you can see, multiple providers are giving their items to the system to be recommended to the users. The multi-provider recommender can be seen in two different types depending on whether or not the providers have preferences towards what type of users they want to reach out to:

- **Without Provider Preferences** In this type of multi-provider recommendation the system is aware that there are multiple providers behind the items and they should get a fair exposure in the recommendations. However, the providers themselves have no preferences towards certain users to whom their items want to be recommended. For example, on the Kiva.org microlending platform, the items which are recommended are loans. That means the providers in this case



(a) Heterogeneous multi-receiver recommendation



(b) Homogeneous multi-receiver recommendation

Figure 2: The architectures for a) Heterogeneous and b) Homogeneous multi-receiver recommendations

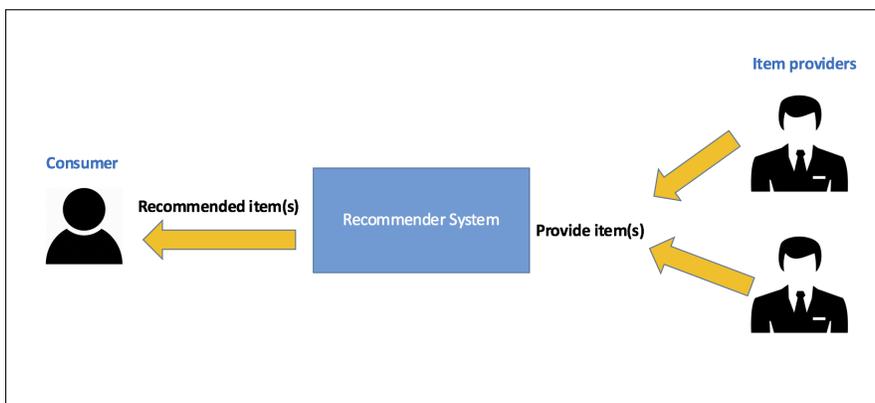


Figure 3: Multi-provider recommender system

are the people who asked for loan on the platform (i.e. the borrowers). In this example, the borrowers do not really care to whom their loan requests are recommended as long as the loan has a good chance of being funded. In other words, they do not have a specific target audience in mind.

- **With Provider Preferences** The providers in this case actually have certain type of audience in mind and the recommender system tries to fulfill the providers' preferences. For example, in *computational advertising* an ad should be recommended to a user if s/he falls within the predefined

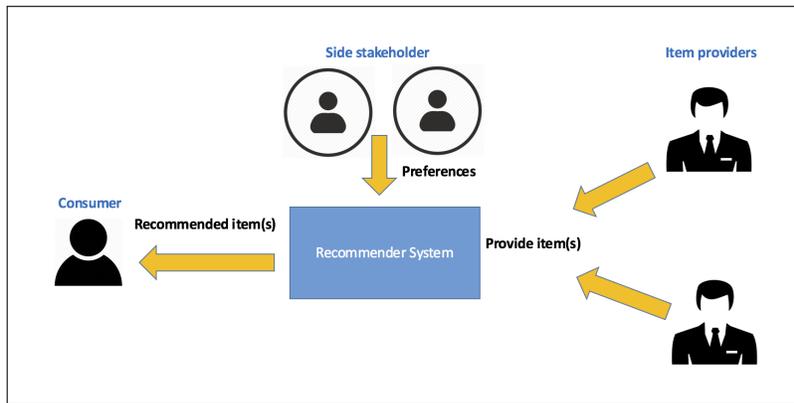


Figure 4: Recommendation with side stakeholders

groups of users who should get that particular ad [26]. In other words, in addition to the ad not being annoying to the user (sadly it is often the case), the user should also be acceptable from the advertiser’s perspective as they might want to reach certain users with a specific characteristics to maximize their ad efficiency.

It is worth mentioning that, in some literature, the multi-provider recommendation is referred to as two-sided market [16] where on one side we have the consumers and on the other side we have the item providers. We believe the name multi-provider recommendation can better explain the behavior of these type of recommendation systems as it clearly emphasizes the other side of the recommendation (besides the consumers): the item providers.

### 3.3 Recommendation with Side Stakeholders

Stakeholders in a recommender system may not be always a direct part of the recommendation interaction. In other words they do not have to necessarily be the consumer of recommendations nor do they have to be the provider of the items being recommended. In some recommendation problems, there are other parties that are being affected by a given recommendation. In such cases, the satisfaction of these *side stakeholders* should be taken into account in recommending items to users. For instance, on Uber Eats, when a list of restaurants is being recommended to a user, the person who is responsible for delivering food to the user is also potentially affected by this recommendation. Not only are the driver’s preferences for a given recommendation important, the overall economy of the amount of work load given to the delivery partners is also something that the system should consider.

The main characteristics of this type of recommendation systems is the fact that other stakeholders who are affected by the recommendations are not the entities who provide the recommendations (multi-provider) nor are they the users who receive the recommendations (multi-receiver). As in the Uber Eats example, the delivery partners are not receiving food recommendations nor are they providing the recommendations (like restaurants) but they still are affected by these recommendations. Figure 4 shows the architecture of a recommendation platform with side stakeholders.

In this model, there are some side stakeholders in addition to the receivers and item providers whose preferences should be taken into account by the recommender system. One interesting challenge in this type of systems is how the system should aggregate those preferences and to what extent each stakeholder’s preferences should be incorporated or prioritized in this decision making process.

**3.3.1 Value-aware recommendation.** A frequent instantiation of a recommendation with side stakeholders is known as value-aware recommendation [5, 18] where the recommendation platform is also considered as a stakeholder as it may have some goals and preferences with respect to the recommendations—considerations such as profit maximization [4, 11], long-tail promotion [3, 21] are examples of such goals and preferences. These systems are called *value-aware recommendation* because the recommender system needs to make sure, in addition to the users’ satisfaction, the recommendations bring some sort of value to the business. For example, authors in [11] proposed a value-aware movie recommendation that takes both the price for each movie and also the user satisfaction into account in generating the recommendations. They showed it was possible to significantly increase the revenue with a negligible loss in accuracy.

### 3.4 Hybrid

These were the most commonly observed classes of multi-stakeholder recommendation. One can imagine a combination of any of these classes. For example, the Uber Eats platform is a combination of multi-provider recommender system (different restaurants) and recommender with side stakeholders (the delivery partners). In fact, Uber Eats could be also a value-aware recommender system as the Uber Eats itself might also have certain business goals in mind.

## 4 FAIRNESS IN MULTI-STAKEHOLDER RECOMMENDATION

Recently, fairness-aware recommendation has attracted a lot of attention from the recommender systems research community [6, 13, 24]. As we discussed earlier, there is a very close connection between multi-stakeholder recommendation and fairness-aware recommendation as having multiple stakeholders gives rise to the

questions of fairness in recommendation. As in [1, 8] we can see different classes of fairness, distinguished by the fairness issues that arise relative to what authors call different sides of the platform: consumers (C-fairness) and providers (P-fairness). In this paper, as we saw in section 3.4, there could be some stakeholders that are neither the receiver nor are they the provider of the recommendations. Therefore, in addition to the mentioned fairness types, we define another type which accounts for the fairness towards other affected stakeholders in the system. We call this type of fairness *S-Fairness* which refers to the idea of the recommender system being fair to the *side* stakeholders who are not directly participating in the recommendation process but rather being affected by such transaction. Therefore, in a multi-stakeholder recommendation platform we can observe the following types of fairness:

- **C-fairness:** A recommender system distinguished by C-fairness is one that must take into account the disparate impact of recommendation on protected classes of recommendation consumers. For example, in a job recommender system that has a C-fairness requirement, the recommendations should be fair towards the users in the protected class (as defined by gender, age, nationality, etc.) relative to other users. In this paper, however, we extend the *C-fairness* to any fairness concerns that the system might have with respect to the consumers of the recommendation. For instance, in a group recommender system the *C-fairness* could refer to taking into account the preferences of different members of the group in a fair way and avoid ignoring certain users' preferences for a long term [23].
- **P-fairness:** A recommender system that has a P-fairness requirement should treat the providers of the items in a fair way. For example, in job recommendation platform, it could mean that minority-owned businesses have their jobs recommended to qualified candidates.
- **S-Fairness** There are examples of multi-stakeholder recommendation in which some stakeholder are not the receiver nor are they the providers of the recommendations. Instead, they act as *side* stakeholders whose needs and preferences should be taken into account in a fair way. For example, in the Uber Eats example, the delivery partners are side stakeholders and, therefore, the system needs to generate the recommendations in a way that the fairness considerations towards these delivery partners such as fair work load, fair commute distance etc. are being addressed properly.

In particular applications, multiple fairness concerns may arise on different sides of the interaction. Thus, a system may have any combination of these fairness considerations in play at once: CP-fairness, for both consumers and providers, as noted in [1], but also any of the others.

## 5 CONCLUSION AND FUTURE WORK

In this paper we defined several most commonly observed classes of multi-stakeholder recommendation. We argued how some of the already existing recommendation problems such as group recommendation and reciprocal recommendation are also examples of multi-stakeholder recommendation. In addition, we discussed

several different types of fairness that are important to address in a multi-stakeholder recommendation. For future work, we intend to develop algorithmic solutions for each of these different types of multi-stakeholder recommendation.

## REFERENCES

- [1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2019. Beyond Personalization: Research Directions in Multistakeholder Recommendation. *arXiv preprint arXiv:1905.01986* (2019).
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Recommender systems as multi-stakeholder environments. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP2017)*. ACM.
- [3] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing Popularity Bias in Recommender Systems with Personalized Re-ranking. In *Florida AI Research Symposium (FLAIRS)*. ACM, To appear.
- [4] Amos Azaria, Avinatan Hassidim, Sarit Kraus, Adi Eshkol, Ofer Weintraub, and Irit Netanel. 2013. Movie recommender system for profit maximization. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 121–128.
- [5] Robin Burke, Gediminas Adomavicius, Ido Guy, Jan Krasnodebski, Luiz Pizzato, Yi Zhang, and Himan Abdollahpouri. 2017. VAMS 2017: Workshop on Value-Aware and Multistakeholder Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 378–379.
- [6] Robin Burke, Nasim Sonboli, Masoud Mansoury, and Aldo Ordoñez-Gauger. 2017. Balanced Neighborhoods for Fairness-aware Collaborative Recommendation. In *Workshop on Responsible Recommendation (FATRec)*.
- [7] Robin D. Burke, Himan Abdollahpouri, Bamshad Mobasher, and Trinadh Gupta. 2016. Towards Multi-Stakeholder Utility Evaluation of Recommender Systems. In *Workshop on Surprise, Opposition, and Obstruction in Adaptive and Personalized Systems, UMAP 2016*.
- [8] Burke, Robin. 2017. Multisided Fairness for Recommendation. In *Workshop on Fairness, Accountability and Transparency in Machine Learning (FATML)*. Halifax, Nova Scotia, 5 pages.
- [9] Oscar Celma. 2010. Music recommendation. In *Music recommendation and discovery*. Springer, 43–85.
- [10] Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalčić. 2018. *Group recommender systems: An introduction*. Springer.
- [11] Dietmar Jannach and Gediminas Adomavicius. 2017. Price and Profit Awareness in Recommender Systems. In *Proceedings of the ACM RecSys 2017 Workshop on Value-Aware and Multi-Stakeholder Recommendation*. Como, Italy.
- [12] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Enhancement of the Neutrality in Recommendation. In *Decisions@ RecSys*. 8–14.
- [13] Eric L. Lee, Jing-Kai Lou, Wei-Ming Chen, Yen-Chi Chen, Shou-De Lin, Yen-Sheng Chiang, and Kuan-Ta Chen. 2014. Fairness-Aware Loan Recommendation for Microfinance Services. ACM Press, 1–4.
- [14] George Lekakos and Petros Caravelas. 2008. A hybrid approach for movie recommendation. *Multimedia tools and applications* 36, 1-2 (2008), 55–70.
- [15] Judith Masthoff. 2011. Group recommender systems: Combining individual models. In *Recommender systems handbook*. Springer, 677–702.
- [16] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. ACM, 2243–2251.
- [17] Ioannis Paparizos, B Barla Cambazoglu, and Aristides Gionis. 2011. Machine learned job recommendation. In *Proceedings of the fifth ACM Conference on Recommender Systems*. ACM, 325–328.
- [18] Changhua Pei, Xinru Yang, Qing Cui, Xiao Lin, Fei Sun, Peng Jiang, Wenwu Ou, and Yongfeng Zhang. 2019. Value-aware Recommendation based on Reinforced Profit Maximization in E-commerce Systems. *arXiv preprint arXiv:1902.00851* (2019).
- [19] Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska, and Judy Kay. 2010. RECON: a reciprocal recommender for online dating. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 207–214.
- [20] Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska, and Judy Kay. 2010. RECON: a reciprocal recommender for online dating. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 207–214.
- [21] Idan Szpektor, Aristides Gionis, and Yoelle Maarek. 2011. Improving recommendation for long-tail queries via templates. In *Proceedings of the 20th international conference on World wide web*. ACM, 47–56.
- [22] Peng Xia, Benyuan Liu, Yizhou Sun, and Cindy Chen. 2015. Reciprocal Recommendation System for Online Dating. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. ACM, 234–241.

- [23] Lin Xiao, Zhang Min, Zhang Yongfeng, Gu Zhaoquan, Liu Yiqun, and Ma Shaoping. 2017. Fairness-aware group recommendation with pareto-efficiency. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 107–115.
- [24] Sirui Yao and Bert Huang. 2017. Beyond parity: Fairness objectives for collaborative filtering. In *Advances in Neural Information Processing Systems*. 2921–2930.
- [25] Hongtao Yu, Chaoran Liu, and Fuzhi Zhang. 2011. Reciprocal recommendation algorithm for the field of recruitment. *JOURNAL OF INFORMATION & COMPUTATIONAL SCIENCE* 8, 16 (2011), 4061–4068.
- [26] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1077–1086.
- [27] Yong Zheng, Nastaran Ghane, and Milad Sabouri. 2019. Personalized Educational Learning with Multi-Stakeholder Optimizations. In *Adjunct Proceedings of the ACM conference on User Modelling, Adaptation and Personalization*. ACM.