Fairness in Job Recommendation under Quantity Constraints

Yunqi Li,[1] Michiharu Yamashita, [2] Hanxiong Chen, [1] Dongwon Lee, [2] Yongfeng Zhang [1]

[1] Department of Computer Science, Rutgers University, New Brunswick, NJ, USA
[2] College of Information Sciences and Technology, Penn State University, University Park, PA, USA
yunqi.li@rutgers.edu, michiharu@psu.edu, hanxiong.chen@rutgers.edu, dongwon@psu.edu, yongfeng.zhang@rutgers.edu

Abstract

Job recommendation platforms such as LinkedIn play an important role in effectively matching job postings and job seekers. The problem of recommending jobs to potential candidates is essentially different from traditional recommendation problems such as recommending movies or music to users since a job posting usually targets for hiring a limited number of employees, which means that quantity constraints are imposed on recommendation conversions. Therefore, recommending a job to an unnecessarily large number of users will not only cost companies excessive efforts in the process of reviewing and interviewing, but also result in intense competition among job seekers on the same job. To solve this problem in a more efficient way, we model the job recommendation task as a resource allocation problem where jobs are recommended under certain quantity constraints. Moreover, to enhance the fairness of users in the job recommendation task, we expect that the recommendation results are consistent to users of different groups after the allocation process. To this end, we propose a framework based on a re-ranking approach with quantity and fairness constraints that generates high-quality recommendations while mitigating the disparity of recommendation results between different user groups. In experiments, we focus on a practical case of user fairness in job recommendation to improve salary fairness over gender. The experimental results on a real-world dataset with various representative recommendation algorithms show that our approach not only generates better distribution of job postings in recommendation lists, but also improves user group fairness while maintaining competitive recommendation qualities.

Introduction

Job recommendation is an important research topic that can benefit both job seekers and job providers (Al-Otaibi and Ykhlef 2012). The task of recommending jobs to potential job seekers is fundamentally different from recommending traditional items such as movies or music to users (Borisyuk, Zhang, and Kenthapadi 2017). In traditional recommendation problems, the quantities of items are usually assumed as infinite. Suppliers expect their items to be recommended to as many users as possible to improve conversions and profits. However, in the scenario of recommending jobs, it is not desirable if the system recommends a job posting to an unnecessarily large number of job seekers. From the job providers side, a job posting usually only intends to hire one or a few employees. Therefore, the workload of reviewing resumes and conducting interviews can greatly exceed expectation if an unnecessarily large number of applications are received. From the job seekers side, a person usually can accept only one job offer. If too many job seekers compete for the same job, the competition for getting the job will become more intense, thus reducing the chances of job seekers being hired. Moreover, modern recommender systems are usually built based on the principle of collaborative filtering, therefore, they easily suffer from popularity bias issues (Abdollahpouri, Burke, and Mobasher 2017, 2019), i.e., the popular items will get more exposure than those less interacted ones. In job recommendation scenario, job providers may leave the recommendation platform if they always get low exposure and receive too few applications when their postings expire, which may further reduce the chance for job seekers to find satisfactory jobs. Therefore, it is also important to ensure a reasonable amount of exposure for job postings to facilitate the long-term development of job recommendation platforms.

Based on aforementioned considerations, we propose a novel approach to model the job recommendation task as a resource allocation problem. Online resource allocation aims to divide an online resource from producers to consumers with their preferences over the resource. The allocation process is usually implicitly realized through search and recommendation since enforced service allocation is not granted by law (Zhang, Zhang, and Friedman 2017). In the job recommendation problem, we treat jobs as resources with limited quantities and aim to recommend each job posting to the most qualified job seekers under certain quantity restrictions. Specifically, each job is assigned with an upper bound for the allocation quantity according to their hiring demands to avoid overwhelmed exposure in the recruitment process. Meanwhile, each job is also assigned with a lower bound for the allocation quantity to ensure a reasonable minimum exposure for job postings. Moreover, the fairness demands of job seekers also need to be considered during the allocation process. In this paper, we consider group fairness for job seekers and aim to allocate/recommend job postings equally across different user groups.

Technically, we introduce a framework to achieve fair resource allocation for job recommendation. The framework is designed based on a post-processing approach with the advantage of making no assumption on the underlying recommendation model and offer model-agnostic flexibility. Specifically, we provide a re-ranking method with allocation quantity constraints and group fairness constraints to re-rank the recommendation lists generated from a base recommender algorithm. We conduct experiments on a real-world dataset with several representative recommendation models to show the effectiveness of our method. In experiments, we study the unfair performance of base models which tend to recommend more lower-salary jobs to female job seekers, and show that our method can generate fairer recommendations by mitigating the salary difference of recommended jobs to users of different genders. The experimental results show that our method can reduce unfairness between protected and advantaged user groups and generate a more reasonable distribution of job postings in top-N recommendation lists without losing the recommendation accuracy.

Related Work

As recommender systems have multiple stakeholders, the fairness of recommendation tasks can be put forward from different perspectives (Burke 2017; Pitoura, Stefanidis, and Koutrika 2021). Many fairness-related works concern unfairness issues from the item side, and focus on the popularity bias in recommendations. Such problem can be addressed by increasing the number of unpopular items, or otherwise, the overall catalog coverage in the final recommendation list (Adomavicius and Kwon 2011; Kamishima et al. 2014; Abdollahpouri, Burke, and Mobasher 2017, 2019). Some other works also consider user-side unfairness issues, for example, Lin et al. (Lin et al. 2017) consider user fairness in group recommendation and propose the optimization framework based on Pareto Efficiency; Machado and Stefanidis(Machado and Stefanidis 2019) focus on the problem of allocating the best members in a fair way between the teams in team recommendations; Li et al. (Li et al. 2021a) study group fairness for users and require recommender system to treat active and inactive users consistently; Li et al. (Li et al. 2021b) achieve personalized counterfactual fairness for users in recommendation. There are also works considering fairness in recommendation from a multi-sided view. Examples include Mehrotra et al. (Mehrotra et al. 2018), which jointly optimizes fairness and performance in two-sided marketplace platforms; Abdollahpouri and Burke (Abdollahpouri and Burke 2019), which introduces several group fairness properties under multi-stakeholder scenario; and Patro et al. (Patro et al. 2020), which explores individual fairness in two-sided platforms from the view of long-term sustainability. In this paper, we consider the particularity of job recommendation tasks and improve user group fairness together with minimum exposure of job postings through fair resource allocation.

Method

In this section, we present the details of our framework. We first introduce the preliminaries and notations about resource allocation and group fairness in job recommendations. After that, we propose a fair resource allocation algorithm based on a re-ranking method under constraints.

Resource Allocation

In the problem of job recommendation, we have user (job seeker) set $\mathcal{U} = \{u_1, u_2, \cdots, u_n\}$, and item (job posting) set $\mathcal{V} = \{v_1, v_2, \cdots, v_m\}$, where $n = |\mathcal{U}|$, $m = |\mathcal{V}|$. The task of recommender system is to provide each user u_i a top-N recommendation list $\{v_1, v_2, \cdots, v_N | u_i\}$ through predicting the matching degree between the user and the candidate items. As we discussed before, we model the task of job recommendation as a resource allocation problem. We consider the quantity of each job posting is limited. Let $\mathbf{q} = [q_1, q_2, \cdots, q_m]$ be the resource quantity vector, where $q_j \geq 0$ is the total amount that job posting v_j can be recommended in the system. The resource allocation problem aims to find an allocation matrix $\mathbf{Q} = [Q_{ij}]_{n \times m}$, where $Q_{ij} \geq 0$ is the quantity that job seeker u_i is recommended with job posting v_j. To meet the resource quantity restriction represented by vector $\mathbf{q}$, we need $\sum_{i=1}^{n} Q_{ij} \leq q_j$ for each job posting v_j. In job recommendation, it is reasonable to require that $Q_{ij} \in \{0, 1\}$, since we usually only consider if or not to recommend a job to a job seeker, and it makes no sense if a job posting appears in a user's top-N recommendation list more than once. Specifically, if Q_{ij} is learned to be 1, the job posting v_j will be recommended to job seeker u_i, otherwise, if $Q_{ij} = 0$, we will not recommend v_j to u_i. To generate top-N recommendation lists for users, we require $\sum_{j=1}^{m} Q_{ij} = N$ for every user u_i.

User Group Fairness

To study group fairness for users, we first divide users into advantaged and disadvantaged (protected) groups based on a certain grouping method. For example, we can divide users based on their sensitive features such as gender to study whether there exists algorithmic bias on gender. Group fairness requires that the protected group should be treated similarly to the advantaged group (Pedreschi, Ruggieri, and Turini 2009). In this paper, we consider dividing users into groups Z_1 and Z_2 in accordance with their sensitive features so that $Z_1 \cap Z_2 = \varnothing$. The notation $\mathcal{F}$ is a metric that can evaluate the recommendation results for users. For example, in job recommendation tasks, users can be concerned about whether they are treated unfairly by always being recommended with jobs of low salaries. In such a case, we can use $\mathcal{F}$ to represent the average salary of jobs in recommendation lists and study whether there is a significant difference on $\mathcal{F}$ between different user groups. We use $\mathcal{F}(\mathrm{Q_i})$ to represent the recommendation result for user i with recommendation list $\mathrm{Q_i}$. The user group fairness in recommendation is defined as follows:

Definition 1 *User Group Fairness (UGF)*

$$\mathbb{E}[\mathcal{F}(\mathbf{Q}) | Z = Z_1] = \mathbb{E}[\mathcal{F}(\mathbf{Q}) | Z = Z_2] \tag{1}$$

The UGF requires that a recommender system offers same-quality recommendation results in expectation for different groups of users. Furthermore, we define ε-fairness recommendation algorithm as follows to measure the user group unfairness of an algorithm:

Definition 2 *(ε-fairness) A recommendation algorithm satisfies ε-fairness if:*

$$UGF\left(Z_1, Z_2, \mathbf{Q}\right)=\left|\frac{1}{\left|Z_1\right|} \sum_{i \in Z_1} \mathcal{F}\left(\mathbf{Q_i}\right)-\frac{1}{\left|Z_2\right|} \sum_{i \in Z_2} \mathcal{F}\left(\mathbf{Q_i}\right)\right| \leq \varepsilon. \tag{2}$$

Here the ε represents the strictness of fairness requirements. It trades off the fairness and the recommendation performance. A smaller ε usually means a stricter fairness requirement and more sacrifice of recommendation accuracy to satisfy fairness demands.

Fair Resource Allocation Algorithm

In this section, we propose a framework to achieve fair resource allocation for job recommendation. To provide model-agnostic flexibility, we propose to re-rank the recommendation lists produced by traditional fairness-unaware recommender systems under quantity and fairness constraints.

For each user-item pair (i, j), a recommender system usually learns a matching score S_{ij} to represent the preference of user u_i towards item v_j. Here we follow the matching scores calculated by the base recommendation model and re-rank the recommendation list so as to achieve the optimal balance between recommendation utility and fairness. To this end, we model users' satisfaction with the recommendation results through a utility function, and we adopt the frequently used personalized utility function (Zhang et al. 2016) which satisfies the diminishing marginal utility requirement:

$$U_{ij}\left(Q_{ij}\right)=\frac{1}{1+e^{-S_{ij}}} \cdot Q_{ij} \tag{3}$$

where $U_{ij}\left(Q_{ij}\right)$ represents the utility when supplying a quantity Q_{ij} of item v_j to user u_i. The intuition of equation (3) is straightforward: the utility is modeled by the sigmoid function so that the utility value U_{ij} is between 0 and 1. Given a user u_i, for items with $Q_{ij}=1$, a larger preference score S_{ij} leads to a larger utility score U_{ij}, which represents higher satisfaction. And for items with $Q_{ij}=0$, the utility score U_{ij} will be 0 since those items are not assigned to the user.

We develop a re-ranking algorithm to maximize the total utility of users under the quantity and fairness constraints to generate fair top-N recommendation lists. The optimization procedure of fair resource allocation for job recommendation is as follows:

$$\begin{gathered}
\max_{Q_{ij}} \quad \sum_{i=1}^{n} \sum_{j=1}^{m} \frac{1}{1+e^{-S_{ij}}} \cdot Q_{ij} \\
\text{s.t.} \quad \operatorname{UGF}\left(Z_1, Z_2, \mathbf{Q}\right)<\varepsilon \\
a_j \leq \sum_{i=1}^{n} Q_{ij} \leq q_j, \quad \sum_{j=1}^{m} Q_{ij}=N, \quad Q_{ij} \in\{0,1\}, \forall i, j
\end{gathered} \tag{4}$$

This objective function can be interpreted as that, for each user, we select N items out of the candidates set and recommend them to the user so that the total utility can be maximized. Meanwhile, these selected items make sure that the top-N recommendation list satisfies the quantity and fairness constraints. Here a_j is the lower bound of the quantity constraint, which requires a minimum exposure for each item. In practice, the value of a_j and q_j can be directly specified by job providers, for example, some job providers require a minimum exposure rate when signing contract with the job recommendation platform. If necessary, the quantity restrictions can also be decided by taking into account the job provider's expected number of hires and the conversion rate of the recommendation platform, for example, the value of a_j and q_j can be set proportional to the number of employees that the job provider plans to hire. The optimization problem here can be solved as a 0–1 integer programming problem. We can find feasible solutions to this problem through fast heuristics[1]. For each user, we rank those N items whose $Q_{ij}=1$ according to the preference score S_{ij} to construct the final recommendation list.

Experiments

In this section, we first briefly describe the dataset, baselines, and experimental settings used for experiments. After that, we evaluate our proposed fair resource allocation framework on top of the baselines to show its desirable performance. To study user group fairness, we consider the important gender-salary fairness in job recommendation, where the goal is to reduce the gender-based wage gap between male and female job seekers without hurting the recommendation accuracy. Specifically, we divide users according to their gender into the male group and the female group. We use $\mathcal{F}\left(\mathbf{Q}_i\right)$ to represent the average salary of user i's recommended jobs in the top-N list, and calculate UGF@N to show the salary gap between the recommended jobs of male and female groups.

Dataset

We use a real-world dataset from a private entity to study fairness in job recommendation. The dataset includes user gender information and company features such as company size and salary (in U.S. dollars). The dataset includes 17,072 user-company interactions, 3,000 users, and 5,105 companies. In our experiments, we randomly split the dataset into the train (80%), validation (10%) and test sets (10%), and all models are trained, validated and tested on the same dataset.

Baselines

Our proposed method is a framework that can be applied on any recommendation model as long as it predicts the preference score for user-item pairs. To evaluate the effectiveness of our framework, we apply it over both shallow and deep recommendation models, including two shallow models (**PMF** and **BiasedMF**), one deep model (**NeuMF**), as well as one sequential model (**STAMP**).

- **PMF** (Mnih and Salakhutdinov 2008): It is a probabilistic matrix factorization algorithm which adds Gaussian prior into the user and item representation distribution.

[1]We use gurobi solver in our experiment. https://www.gurobi.com

Table 1: Recommendation and fairness performance of our method (re-Ranked), and baseline models. The evaluation metrics are calculated based on top-10 predictions in test set. UGF@10 is the difference of average salary between male and female groups.

Model		NDCG@10	HR@10	F1@10	UGF@10
PMF	Baseline	0.0928	0.1295	0.0286	$14,368
	re-Ranked	0.0924	0.1265	0.0280	$967
BiasedMF	Baseline	0.0760	0.1066	0.0236	$16,050
	re-Ranked	0.0727	0.1043	0.0231	$960
NeuMF	Baseline	0.0747	0.1037	0.0230	$10,666
	re-Ranked	0.0740	0.1025	0.0227	$980
STAMP	Baseline	0.0698	0.0808	0.0172	$4,155
	re-Ranked	0.0728	0.0902	0.0190	$952

- **BiasedMF** (Koren, Bell, and Volinsky 2009): It is a matrix factorization algorithm which takes user, item and global bias terms into consideration.
- **NeuMF** (He et al. 2017): It is an algorithm which applies deep neural network with non-linear activation functions to learn matching function for user-item pair.
- **STAMP** (Liu et al. 2018): It is a session-based recommendation model built on attention mechanism to capture user's long-term and short-term preferences.

Evaluation

We use standard ranking metrics Hit Rate (HR), F_1 score, and Normalized Discounted Cumulative Gain (NDCG) to evaluate the top-K recommendation quality and use UGF to evaluate user fairness. We treat the entire non-interacted item set as a user's candidate list (all-item ranking), and compute the metric scores over this candidate list to evaluate the top-N recommendation performance. The results of all metrics in our experiments are averaged over all users.

Implementation Details

We apply Bayesian Personalized Ranking (BPR) (Rendle et al. 2012) loss for all the models. Specifically, we randomly sample one item that a user has never interacted with as the negative sample for each observed user-item pair during each training epoch. We tune the hyper-parameters of all the models to reach their best performance. The best models are selected based on the performance on the validation set within 100 epochs. Since we do not have information about how many employees the company plans to hire, we use the quantity that is proportional to the company size to simulate the hiring demands with the reasonable assumption that larger companies tend to have more opening positions. To simulate the exposure demands of job providers, we set the upper bound of the quantity constraint to 40-60% of the company size and set the lower bound of the quantity constraint as 3 to 5. For the user-side fairness constraint, we set the maximum difference of the average salary between male and female user groups to 1,000 U.S. dollars.

Experimental Results

Recommendation Performance We first compare the recommendation performance of our method with baseline models, and the results are shown in Table.1. We can see that our method achieves competitive recommendation accuracy compared with baseline models. For the sequential model STAMP, our method even gets better results on all the evaluation metrics. These results show that although our method aims to achieve multiple goals including user fairness and more reasonable exposure of job postings, we do not have to sacrifice recommendation accuracy.

Effect of Fairness Constraint We then analyze the fairness improvement and the effect of fairness constraint in the learning process. From the results of UGF@10 in Table.1, we can see that our method is very effective at improving fairness for users. For all the baseline models, we observe that male users tend to get job recommendations of higher salaries than female users. The significant difference in the average salary of the top-10 recommendation lists between male and female groups implies that the fairness-unaware recommendation models tend to treat female users unfairly by recommending more low-salary jobs to them. Moreover, we can see that our method can narrow the gap significantly, while maintaining the recommendation quality. The fairer recommendations will provide female users more opportunities to get high-salary jobs and better careers. To study the effect of fairness constraint, we also conduct experiments of re-ranking the recommendation lists based only on fairness constraint. We show the recommendation performance of NeuMF as an example in Fig.(1a), and other models have similar trends. We find that adding the fairness constraint will reduce the recommendation quality, which is reasonable since it narrows the solution space of the learning problem.

Effect of Quantity Constraint We also study the effect of quantity constraints in the re-ranking process. Firstly, we conduct experiments of doing re-ranking based only on quantity constraints. From the results of Fig.(1a), we observe that adding quantity constraints can improve the recommendation performance since the quantity constraints help the recommendation model converge to a quantitatively reasonable solution without over-exposing some jobs or under-exposing others. Besides, we group companies according to their hiring scale, and plot the average exposure of company groups in the recommendation results in Fig.(1b). The groups $G0$, $G1$, $G2$ are companies with size less than 1000, 1000 to 5000, and larger than 5000, respectively. From the figure, we can see that for the companies that have relatively smaller sizes ($G0$ and $G1$) and only plan to hire a few employees, our method can improve their exposure in recommendation platforms. Such companies are often less popular in the marketplaces and suffer from low exposure and fewer applications. Therefore, improving their exposure can help platforms to retain small companies and to improve the opportunities for users to find satisfactory jobs. For those companies ($G2$) of larger size and more popular in the platform, our method is effective to help them get rid of unnecessary applications since our method reduces exposure for them without hurting the recommendation accuracy. There-

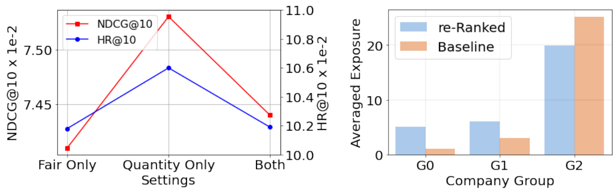

(a) (b)

Figure 1: (a) Recommendation performance of re-ranking on NeuMF under different constraints; (b) Averaged exposure of job postings w.r.t different company groups.

fore, the results show that our method can generate fairer and more accurate job recommendations, help platforms to improve user satisfaction, and may also contribute to the long-term prospering of the job recommendation platform and job markets.

Conclusion

In this paper, we emphasize the importance and necessity to consider the quantity limitations and fairness demands in job recommendation task, and propose to model the job recommendation as a resource allocation problem. We propose a framework to achieve fair resource allocation for job recommendations through a post-processing method. Experimental results show that our method is able to improve the recommendation fairness between protected and advantaged user groups while maintaining desirable recommendation accuracy.

References

Abdollahpouri, H.; and Burke, R. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158*.

Abdollahpouri, H.; Burke, R.; and Mobasher, B. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, 42–46.

Abdollahpouri, H.; Burke, R.; and Mobasher, B. 2019. Managing popularity bias in recommender systems with personalized re-ranking. *arXiv preprint arXiv:1901.07555*.

Adomavicius, G.; and Kwon, Y. 2011. Improving aggregate recommendation diversity using ranking-based techniques. *IEEE Transactions on Knowledge and Data Engineering*, 24(5): 896–911.

Al-Otaibi, S. T.; and Ykhlef, M. 2012. A survey of job recommender systems. *International Journal of Physical Sciences*, 7(29): 5127–5142.

Borisyuk, F.; Zhang, L.; and Kenthapadi, K. 2017. LiJAR: A system for job application redistribution towards efficient career marketplace. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1397–1406.

Burke, R. 2017. Multisided fairness for recommendation. *arXiv preprint arXiv:1707.00093*.

He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural Collaborative Filtering. In *WWW*, 173–182.

Kamishima, T.; Akaho, S.; Asoh, H.; and Sakuma, J. 2014. Correcting Popularity Bias by Enhancing Recommendation Neutrality. In *RecSys Posters*.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*, (8): 30–37.

Li, Y.; Chen, H.; Fu, Z.; Ge, Y.; and Zhang, Y. 2021a. User-oriented fairness in recommendation. In *Proceedings of the Web Conference 2021*, 624–632.

Li, Y.; Chen, H.; Xu, S.; Ge, Y.; and Zhang, Y. 2021b. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1054–1063.

Lin, X.; Zhang, M.; Zhang, Y.; Gu, Z.; Liu, Y.; and Ma, S. 2017. Fairness-aware group recommendation with pareto-efficiency. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, 107–115.

Liu, Q.; Zeng, Y.; Mokhosi, R.; and Zhang, H. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *SIGKDD*, 1831–1839. ACM.

Machado, L.; and Stefanidis, K. 2019. Fair team recommendations for multidisciplinary projects. In *2019 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, 293–297. IEEE.

Mehrotra, R.; McInerney, J.; Bouchard, H.; Lalmas, M.; and Diaz, F. 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*, 2243–2251.

Mnih, A.; and Salakhutdinov, R. R. 2008. Probabilistic matrix factorization. In *Advances in neural information processing systems*, 1257–1264.

Patro, G. K.; Biswas, A.; Ganguly, N.; Gummadi, K. P.; and Chakraborty, A. 2020. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. In *Proceedings of The Web Conference 2020*, 1194–1204.

Pedreschi, D.; Ruggieri, S.; and Turini, F. 2009. Measuring discrimination in socially-sensitive decision records. In *Proceedings of the 2009 SIAM international conference on data mining*, 581–592. SIAM.

Pitoura, E.; Stefanidis, K.; and Koutrika, G. 2021. Fairness in rankings and recommendations: An overview. *The VLDB Journal*, 1–28.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.

Zhang, Y.; Zhang, Y.; and Friedman, D. 2017. Economic recommendation based on pareto efficient resource allocation. Technical report, WZB Discussion Paper.

Zhang, Y.; Zhao, Q.; Zhang, Y.; Friedman, D.; Zhang, M.; Liu, Y.; and Ma, S. 2016. Economic recommendation with surplus maximization. In *Proceedings of the 25th International Conference on World Wide Web*, 73–83.