

A Survey on Multisided Fairness in Recommender Systems

Reza Shafiloo^[0009–0002–4827–2112], Maria Stratigi^[0000–0003–2482–4605], and
Kostas Stefanidis^[0000–0003–1317–8062]

Data Science Research Centre, Tampere University, Finland
`{reza.shafiloo,maria.stratigi,konstantinos.stefanidis}@tuni.fi`

Abstract. While traditional recommender systems primarily optimize recommendation accuracy, growing attention has been devoted to fairness, as decisions can create unequal outcomes for different affected parties. This challenge is particularly important in multisided recommender systems, where recommendations influence not only consumers but also providers, items, and platforms. In this survey, we review the literature on multisided fairness in recommender systems. We discuss key stakeholder groups, fairness objectives, and fairness-aware recommendation methods, including pre-processing, in-processing, and post-processing approaches. We further examine recommendation settings with limited item availability, where recommendation becomes an allocation problem and fairness must be considered across multiple interactions. Finally, we outline open challenges and future research directions.

Keywords: Multisided Fairness · Multi-Stakeholder Recommendations · Fairness-aware Recommendations.

1 Introduction

Recommender systems increasingly mediate access to information, products, services, and opportunities across digital platforms, including e-commerce, news, health, education, tourism, and entertainment systems [28, 9]. They typically learn from historical interactions, such as ratings, clicks, reviews, and purchases, and use these signals to predict which items are likely to match each user’s interests. However, because recommendation models are trained on historical data, they may reproduce and amplify existing biases. For example, popular items may receive repeated exposure, while less popular items remain less visible [15]. Similarly, active users may receive higher-quality recommendations because they provide more interaction data, whereas less active users may receive lower-quality recommendations. Fairness in recommendation becomes particularly challenging in multisided environments, where recommendation decisions affect several stakeholders simultaneously. Consumers may expect accurate and useful recommendations, providers may seek fair visibility for their items, items may require balanced exposure, and platforms may need to balance engagement, revenue,

and long-term ecosystem health. These objectives are not always aligned: improving fairness or utility for one stakeholder may reduce the opportunities or benefits of another [28, 21]. Therefore, fairness-aware recommendation should not be studied only from a single stakeholder perspective.

These challenges have motivated growing research on multisided fairness in recommender systems. Multisided fairness examines how recommendation decisions distribute utility, exposure, opportunities, and benefits across multiple affected parties, and how potentially competing fairness objectives can be balanced. A particularly underexplored setting arises when recommended items have limited availability, such as limited seats, tickets, appointments, or products. In such cases, recommendation is not only a ranking problem but also an allocation problem, because assigning a scarce item to one user can reduce the opportunity for other users to receive the same item. In this survey, we review the emerging literature on multisided fairness in recommender systems. We discuss key stakeholder groups, fairness objectives, and fairness-aware recommendation methods. We further examine recommendation settings with limited item availability, where fairness must be considered across multiple interactions and over time. Finally, we outline open challenges and future research directions.

2 Fairness in Recommender Systems

Bias and Fairness Concerns. Recommender systems learn from historical user interactions, such as ratings, clicks, purchases, and reviews. As a result, they can inherit and amplify biases present in the data or introduced by the recommendation process itself. One of the most widely studied examples is popularity bias, where popular items receive disproportionately high exposure while less popular items remain under-represented. Since popular items typically generate more interaction data, recommendation models can estimate their relevance more accurately, which further reinforces their visibility. This feedback loop can limit catalog coverage, reduce recommendation diversity, and hinder the discovery of niche or long-tail items [17]. Biases can also affect different entities involved in recommendation. Consumers with rich interaction histories often receive more accurate recommendations than consumers with limited historical data, creating disparities in recommendation quality [10]. Similarly, providers may experience unequal visibility when recommendation algorithms repeatedly favor well-established providers over smaller or less active ones. Such disparities illustrate that recommendation quality alone is insufficient for evaluating a recommender system. The distribution of opportunities, exposure, and benefits across affected entities must also be considered.

Fairness has therefore emerged as an important objective in recommender systems. It refers to the equitable treatment of entities affected by recommendation decisions [23]. Existing work studies fairness from different perspectives, including consumer fairness, provider fairness and item fairness [10]. Depending on the application, fairness may aim to ensure comparable recommendation quality across consumers, balanced exposure among providers, equitable visibil-

ity of items, or reduced disparities between protected and non-protected groups. Since these objectives are often conflicting, fairness is typically formulated as a trade-off between recommendation utility and equitable outcomes.

Fairness-Aware Recommendation Methods. Existing fairness-aware recommendation methods are commonly grouped into pre-processing, in-processing, and post-processing approaches [23]. Pre-processing methods address unfairness before model training by modifying the input data. Their objective is to reduce biases embedded in historical interactions, ratings, or user behavior patterns. Representative approaches include data balancing, causal data repair, synthetic data generation, and bias mitigation techniques that alter the training distribution before learning takes place [26, 7, 5, 13]. By improving the quality and representativeness of the input data, these methods aim to prevent unfair patterns from being learned by the recommendation model. In-processing methods incorporate fairness directly into the learning process. These approaches modify the recommendation algorithm by introducing fairness-aware objective functions, regularization terms, constraints, or optimization strategies [4, 38, 32, 36]. The model is trained to jointly optimize recommendation quality and fairness objectives, allowing fairness considerations to become part of the model behavior itself. As a result, in-processing methods can provide stronger fairness guarantees, although they often require changes to the underlying recommendation architecture. Post-processing methods introduce fairness after recommendation scores or ranked lists have been produced. These methods adjust the output of an existing recommender through re-ranking, score modification, or allocation strategies [31, 12, 8]. Because they do not require retraining the recommendation model, post-processing approaches are particularly attractive in deployed systems where modifying the underlying algorithm may be costly or impractical. Existing work further distinguishes between non-parametric methods, which rely on heuristic or rule-based adjustments, and parametric methods, which learn parameters that guide the fairness-aware re-ranking process [18].

Overall, fairness-aware recommendation seeks to balance recommendation utility with equitable outcomes across affected entities. While early research primarily focused on consumer-side fairness, recent studies increasingly recognize that recommendation decisions affect multiple stakeholders simultaneously. This observation has motivated the development of multisided fairness approaches, which consider the interests and fairness objectives of several stakeholder groups within the same recommendation process.

3 Multisided Fairness in Recommender Systems

Motivation. Traditional fairness-aware recommender systems primarily focus on the fairness experienced by consumers. However, recommendation decisions often affect multiple entities simultaneously. In many application domains, recommendations influence not only the users who receive them, but also the providers that supply items, the items competing for visibility, and the platforms that mediate these interactions. Consequently, fairness cannot be viewed

solely as a consumer-side objective. This observation has motivated the study of multisided fairness in recommender systems. A recommender system is considered multisided when recommendation outcomes affect the utility, opportunities, or benefits of multiple stakeholder groups. Typical examples include online marketplaces, media platforms, recruitment systems, and crowdsourcing platforms, where recommendations directly influence both demand-side and supply-side participants [19, 16, 22]. In these environments, recommendation decisions determine not only which items consumers discover, but also which providers gain visibility and which items receive exposure. The need for multisided fairness arises because recommendation objectives across stakeholders are often interdependent. Improving recommendation quality for consumers may concentrate exposure among a small number of providers, while increasing provider visibility may reduce consumer utility. Similarly, platform objectives such as revenue maximization or marketplace sustainability may not always align with the interests of consumers or providers [22, 21, 24]. As a result, recommendation decisions frequently involve trade-offs among competing fairness and utility objectives.

The literature increasingly recognizes that fairness should be evaluated beyond a single stakeholder perspective [6, 34]. Existing studies have shown that recommendation algorithms can amplify popularity bias, interaction bias, and exposure inequalities, creating systematic advantages for active consumers, popular items, or dominant providers [29, 30, 32]. If such effects remain unaddressed, they can reduce marketplace diversity, discourage provider participation, and weaken user trust in the recommendation process [22, 34]. Therefore, multisided fairness aims to balance the interests of multiple stakeholders while maintaining the overall effectiveness and sustainability of the recommender system.

Stakeholders. The literature identifies several recurring stakeholder categories in recommender systems. Although the exact stakeholders depend on the application domain, consumers, providers, items, and platforms constitute the most common stakeholder groups.

Consumers. Consumers are the recipients of recommendations and are traditionally the primary focus of recommender systems. Their objective is to receive relevant, useful, and personalized recommendations. Fairness at the consumer level typically concerns the equitable distribution of recommendation quality across users or user groups [22, 21, 35]. For example, in recruitment systems, consumers may correspond to job seekers or candidates [27, 19].

Providers. Providers supply the items or services that are recommended. Their objective is to obtain visibility, engagement, and economic benefit through recommendation exposure. Provider fairness aims to prevent recommendation opportunities from being concentrated among a small subset of providers and seeks a more equitable distribution of exposure across the provider population [37, 34, 24, 30]. Depending on the domain, providers may correspond to merchants, content creators, publishers, organizations, or employers.

Items. Items can also be treated as stakeholders when fairness concerns the allocation of recommendation exposure among individual items. This perspective is particularly relevant in the presence of popularity bias, where a small number

of highly popular items dominate recommendation lists while many relevant items remain under-exposed [29]. Item-level fairness is also important in settings where providers maintain large inventories and fairness is desired both across and within provider catalogs.

Platforms. The platform acts as an intermediary between consumers and providers and controls the recommendation process. In addition to facilitating interactions, platforms often pursue objectives such as revenue generation, user retention, marketplace growth, and long-term sustainability [11, 20, 25, 22]. For this reason, several studies explicitly treat the platform as a stakeholder whose objectives must be balanced against consumer and provider fairness.

Domain-Specific Stakeholders. Some application domains involve additional actors whose interests are influenced by recommendation decisions. For example, recruitment systems may include recruiters and hiring managers [27, 19], while crowdsourcing platforms may involve workers, requesters, and workplaces [16]. These examples highlight that stakeholder definitions should be adapted to the characteristics of the application and the decision-making process supported by the recommender system.

Overall, multisided fairness requires considering how recommendation decisions affect multiple stakeholders simultaneously. Because stakeholder objectives are often competing, fairness is increasingly viewed as a problem of balancing utility, exposure, opportunities, and long-term benefits across all affected parties.

Methods. Similar to fairness-aware recommender systems, methods for multisided fairness can be categorized into pre-processing, in-processing, and post-processing approaches, while the specific techniques used in each category may differ according to the stakeholder, fairness goal, and system setting. In addition, many recent approaches formulate multisided fairness as a multi-objective optimization problem, where the utilities of different stakeholders are jointly considered during recommendation generation [11, 33].

Pre-processing methods introduce fairness before model training by modifying the data used to learn recommendations. These methods are useful when unfairness comes from biased interaction data, popularity imbalance, or under-representation of some stakeholder groups. Different pre-processing strategies have been used in the literature. Some approaches inject synthetic ratings for tail items to reduce popularity bias and improve long-tail visibility [14]. Other methods balance the training data by changing the representation of protected groups, such as increasing the presence of women artists in music recommendation [2]. A further strategy uses upsampling to increase the influence of minority providers or under-represented provider groups during relevance learning [3]. In multi-stakeholder settings, these methods can support provider and item fairness by giving less visible providers, minority groups, or long-tail items more learning signals. However, pre-processing may be less effective when unfairness is mainly created during ranking, exposure allocation, or platform-level optimization.

In-processing methods integrate fairness directly into the learning or optimization process. These methods modify the objective function, representation learning, or training weights so that the model learns to balance recommenda-

tion quality with stakeholder fairness. Some approaches use adaptive weighting to increase the effect of disadvantaged users or under-exposed providers during training [35]. Others use adversarial learning, graph neural networks, robust optimization, or debiasing techniques to reduce provider bias, regional bias, or group-level performance gaps [24, 27, 36, 32]. In this group, multi-objective optimization is often used as a learning strategy, because the model must jointly consider consumer utility, provider exposure, item visibility, and sometimes platform revenue [11, 33]. Depending on the application, optimization objectives may combine consumer relevance, provider exposure fairness, item diversity, and platform utility, requiring the system to balance potentially conflicting goals. Thus, in-processing methods are suitable when fairness should be part of the model behavior itself, not only a correction applied after recommendations are produced.

Post-processing methods introduce fairness after a baseline recommender has produced relevance scores or ranked lists. These methods usually re-rank items, adjust scores, or allocate recommendation slots by combining relevance with fairness terms for consumers, providers, or items. They are common in multi-stakeholder recommendation because they can be applied to existing systems without retraining the base model. Some methods balance consumer and producer fairness through optimization-based re-ranking [21, 25]. Others use fair allocation principles, max-min fairness, or exposure guarantees to improve provider visibility while preserving consumer utility. These approaches often view recommendation exposure as a resource that should be distributed across stakeholders according to fairness criteria, such as proportionality, equality of opportunity, or minimum exposure guarantees [22, 34, 37]. Post-processing can also support item and platform fairness through weighted re-ranking or profit-aware re-ranking [29, 20]. In limited item settings, this group can become an allocation method, where scarce items are assigned by fairness-adjusted scores and delivery histories [30]. In such scenarios, fairness must be evaluated not only at the level of individual recommendations but also across repeated allocation decisions over time.

4 Evaluation Metrics

Evaluation in multi-stakeholder recommenders should measure how recommendation outcomes affect different stakeholders and how fairness interventions influence the trade-off between utility and equitable outcomes. It should not only report ranking quality for consumers but also show whether consumers, providers, items, and platforms receive balanced outcomes. Therefore, evaluation metrics can be organized into consumer-side, provider-side, item-side, platform-side.

Consumer-side metrics evaluate whether users receive fair and useful recommendations. Many works measure consumer fairness by comparing recommendation quality across user groups, such as active and inactive users, protected and unprotected groups, or mainstream and non-mainstream users [29, 25, 36, 32]. Common metrics include performance gap, NDCG difference, prediction error gap, and deviation from consumer fairness. Some works also use individual-level metrics, such as the min-max ratio, to check whether some users receive

much worse recommendation quality than others [37]. In hiring recommendation, fairness can also be measured through qualification-aware or opportunity-based metrics [27, 19]. Also, entropy-based metrics can be used to evaluate how evenly consumers receive limited items from different provider categories over time [30].

Provider-side metrics evaluate whether providers receive fair visibility or exposure in recommendation lists. These metrics are important because exposure is linked to economic opportunity and long-term participation. Common metrics include exposure disparity, exposure ratio, disparate visibility, Gini index, entropy, max-min fairness, and variance or coefficient of variation of provider exposure [30, 29, 34, 37, 11, 35]. Some works also use exposure guarantees, such as maximin share, or measure the fraction of providers that receive enough exposure [22]. **Item-side** metrics evaluate whether items receive balanced recommendation opportunities. This is needed because provider fairness does not always imply item fairness. A provider may receive enough exposure, while only a few popular items in its inventory are repeatedly recommended. Common item-side metrics include long-tail coverage, catalog coverage, novelty, entropy of item exposure, average percentage of long-tail items, and deviation from item fairness [29, 14]. Some works also compare exposure across item groups, such as high-price and low-price items, or popular and less popular items [20, 21].

Platform-side and system-level metrics evaluate whether fairness is achieved while preserving broader platform goals. These goals may include revenue, profit, engagement, or long-term sustainability. Some works report revenue-related metrics, such as gains, price, spending, or pay-per-click, together with consumer and provider fairness metrics [11, 20]. In limited item recommendation, system-level fairness can also be measured through allocation-based metrics, such as cumulative delivery distribution [30]. This metric shows whether scarce items reach a broader set of consumers, which can increase the chance of purchase and support platform revenue over time. Overall, evaluating multisided fairness remains challenging because improvements for one stakeholder may come at the expense of another. As a result, recent studies increasingly report multiple fairness and utility metrics simultaneously, providing a more comprehensive view of recommendation quality across the entire recommendation ecosystem [29, 6].

5 Multisided Fairness Under Limited Item Availability

Most fairness-aware recommender systems assume that recommended items are always available. However, many real-world applications involve limited resources, such as event tickets, hotel rooms, healthcare appointments, course seats, or limited-edition products. In these settings, recommendation is not only a ranking problem but also an allocation problem. Specifically, allocating a scarce item to one consumer may reduce or eliminate the opportunity for other consumers to receive the same item. Consequently, fairness depends not only on recommendation relevance but also on how opportunities are distributed among stakeholders over time. This setting introduces additional challenges for multisided fairness, as recommendation decisions affect both consumers competing for limited resources

and providers whose inventories are allocated through the recommendation process. Recent work has begun to address fairness under limited availability. [1] studies recommendation loss in multi-session recommendation settings and propose compensation mechanisms that reduce disparities among consumers who miss access to preferred items. More recently, [30] introduces EQUALI, a framework that extends fairness considerations to both consumers and providers. The framework incorporates historical interactions and previous allocation outcomes into the recommendation process, aiming to distribute limited recommendation opportunities more fairly across multiple stakeholders.

6 Open Challenges and Future Directions

Despite significant progress, several challenges remain open for achieving fair recommendation outcomes across multiple stakeholders.

Balancing Multiple Fairness Objectives. Multisided recommender systems must often balance competing objectives across consumers, providers, items, and platforms. Improving fairness for one stakeholder may negatively affect another stakeholder’s utility or exposure. Designing recommendation algorithms that effectively navigate these trade-offs while maintaining recommendation quality remains a fundamental challenge.

Evaluation Metrics for Multisided Fairness. Existing evaluation metrics primarily focus on individual stakeholders, such as recommendation accuracy for consumers or exposure fairness for providers. However, there is no widely accepted framework for jointly evaluating fairness across multiple stakeholders. Developing evaluation methodologies that capture the overall fairness of a recommendation ecosystem remains an important research direction.

Dynamic and Longitudinal Fairness. Fairness is often evaluated within a single recommendation session, whereas recommendation decisions frequently have long-term consequences. Historical interactions, accumulated exposure, and previous allocation decisions can influence future opportunities for both consumers and providers. This challenge becomes particularly important in settings with limited item availability, where fairness depends on how scarce opportunities are distributed over time.

Explainability and Transparency. As recommender systems increasingly incorporate fairness objectives, stakeholders may seek explanations for recommendation and allocation decisions. Providing transparent and understandable explanations can help users, providers, and platform operators understand potential fairness trade-offs and increase trust in the recommendation process. Integrating explainability with multisided fairness remains largely unexplored.

Real-World Deployment and User Studies. Most existing approaches are evaluated using offline experiments and benchmark datasets. However, the practical impact of multisided fairness interventions on consumer behavior, provider participation, and platform sustainability is still not well understood. More real-world deployments, online evaluations, and user studies are needed to assess the effectiveness of fairness-aware recommendation strategies in practice.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Azzalini, D., Azzalini, F., Criscuolo, C., Dolci, T., Martinenghi, D., Amer-Yahia, S.: Loss compensation in multi-session recommendation under limited availability. In: EDBT. pp. 522–533 (2024)
2. Bauer, C., Ferraro, A.: Strategies for mitigating artist gender bias in music recommendation: a simulation study. In: Music Recommender Systems Workshop (2023)
3. Boratto, L., Fenu, G., Marras, M.: Interplay between upsampling and regularization for provider fairness in recommender systems. *User Modeling and User-Adapted Interaction* **31**(3), 421–455 (2021)
4. Borges, R., Stefanidis, K.: Feature-blind fairness in collaborative filtering recommender systems. *Knowl. Inf. Syst.* **64**(4), 943–962 (2022)
5. Brunet, M.E., Alkalay-Houlihan, C., Anderson, A., Zemel, R.: Understanding the origins of bias in word embeddings. In: Machine Learning. pp. 803–811 (2019)
6. Burke, R., Adomavicius, G., Bogers, T., Di Noia, T., Kowald, D., Neidhardt, J., Özgöbek, , Pera, M.S., Tintarev, N., Ziegler, J.: De-centering the (Traditional) user: Multistakeholder evaluation of recommender systems. *International Journal of Human Computer Studies* **203** (9 2025)
7. Calmon, F., Wei, D., Vinzamuri, B., Natesan Ramamurthy, K., Varshney, K.R.: Optimized pre-processing for discrimination prevention. *Advances in neural information processing systems* **30** (2017)
8. Carraro, D., Bridge, D.: Enhancing recommendation diversity by re-ranking with large language models. *ACM Transactions on Recommender Systems* (2024)
9. Dalla Vecchia, A., Migliorini, S., Quintarelli, E., Stefanidis, K.: Multi-sided fairness in sequential task assignment. *Information Systems* **139**(3), 102700 (7 2026)
10. Deldjoo, Y., Anelli, V.W., Zamani, H., Bellogin, A., Di Noia, T.: A flexible framework for evaluating user and item fairness in recommender systems. *User Modeling and User-Adapted Interaction* pp. 1–55 (2021)
11. Gao, H., Lu, Y., Zhu, N.: A Fairness-Aware Recommender System Based on Genetic Algorithms for User, Provider and Platform Optimization. *ICAACE* (2025)
12. Gómez, E., Contreras, D., Boratto, L., Salamó, M.: Enhancing recommender systems with provider fairness through preference distribution-awareness. *Information Management Data Insights* **5**(1), 100311 (2025)
13. González-Sendino, R., Serrano, E., Bajo, J.: Mitigating bias in artificial intelligence: Fair data generation via causal models for transparent and explainable decision-making. *Future Generation Computer Systems* **155**, 384–401 (2024)
14. Gulsoy, M., Yalcin, E., Bilge, A.: Equirate: balanced rating injection approach for popularity bias mitigation in recommender systems. *PeerJ Computer Science* **11** (2025)
15. Hamedani, E.M., Kaedi, M.: Recommending the long tail items through personalized diversification. *Knowledge-Based Systems* **164**, 348–357 (2019)
16. Huang, W., Li, P., Li, B., Liu, Q., Nie, L., Bao, H.: Three-sided online stable task assignment in spatial crowdsourcing. *Information Sciences* **654**, 119878 (1 2024)
17. Jain, A., Singh, P.K., Dhar, J.: Multi-objective item evaluation for diverse as well as novel item recommendations. *Expert Systems with Applications* **139** (2020)

18. Jin, D., Wang, L., Zhang, H., Zheng, Y., Ding, W., Xia, F., Pan, S.: A survey on fairness-aware recommender systems. *CoRR* **abs/2306.00403** (2023)
19. Kaya, M., Bogers, T.: Mapping stakeholder needs to multi-sided fairness in candidate recommendation for algorithmic hiring. In: *RecSys*. p. 257–267 (2025)
20. Liu, S., Sun, J., Deng, X., Wang, H., Liu, W., Zhu, C., Yang, L.T., Wu, C.: Towards platform profit-aware fairness in personalized recommendation. *International Journal of Machine Learning and Cybernetics* **15**(10), 4341–4356 (2024)
21. Naghiaei, M., Rahmani, H.A., Deldjoo, Y.: Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In: *ACM SIGIR* (2022)
22. Patro, G.K., Biswas, A., Ganguly, N., Gummadi, K.P., Chakraborty, A.: FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. *The Web Conference* (2020)
23. Pitoura, E., Stefanidis, K., Koutrika, G.: Fairness in rankings and recommendations: an overview. *VLDB J.* **31**(3), 431–458 (2022)
24. Qi, T., Wu, F., Wu, C., Sun, P., Wu, L., Wang, X., Huang, Y., Xie, X.: ProFairRec: Provider Fairness-aware News Recommendation. *SIGIR* pp. 1164–1173 (7 2022)
25. Rahmani, H.A., Naghiaei, M., Deldjoo, Y.: A Personalized Framework for Consumer and Producer Group Fairness Optimization in Recommender Systems. *ACM Transactions on Recommender Systems* **2**(3), 1–24 (9 2024)
26. Salimi, B., Rodriguez, L., Howe, B., Suci, D.: Interventional fairness: Causal database repair for algorithmic fairness. In: *SIGMOD*. pp. 793–810 (2019)
27. Schellingerhout, R., Barile, F., Tintarev, N.: Okra: An explainable, heterogeneous, multi-stakeholder job recommender system. In: *Advances in IR* (2025)
28. Shafiloo, R., Kaedi, M., Pourmiri, A.: Considering user dynamic preferences for mitigating negative effects of long-tail in recommender systems. *Information Sciences* **669**, 120558 (2024)
29. Shafiloo, R., Stratigi, M., Peltonen, J., Olsson, T., Stefanidis, K.: The many facets of fairness in recommender systems: Consumers, providers and items. *Information Systems* **137**, 102643 (4 2026)
30. Shafiloo, R., Stratigi, M., Peltonen, J., Stefanidis, K.: Multisided fairness under limited item availability in recommender systems. *Information Sciences* **731** (2026)
31. Stratigi, M., Nummenmaa, J., Pitoura, E., Stefanidis, K.: Fair sequential group recommendations. In: *SAC*. pp. 1443–1452 (2020)
32. Tang, J., Shen, S., Wang, Z., Gong, Z., Zhang, J., Chen, X.: When fairness meets bias: a debiased framework for fairness aware top-n recommendation. In: *RecSys* (2023)
33. Wu, H., Ma, C., Mitra, B., Diaz, F., Liu, X.: A Multi-Objective Optimization Framework for Multi-Stakeholder Fairness-Aware Recommendation. *ACM Transactions on Information Systems* **41**(2) (12 2022)
34. Xu, C., Chen, S., Xu, J., Shen, W., Zhang, X., Wang, G., Dong, Z.: P-mmF: Provider max-min fairness re-ranking in recommender system. In: *The Web Conf.* (2023)
35. Xu, L., Lin, Z., Wang, J., Chen, S., Zhao, W.X., Wen, J.R.: Promoting Two-sided Fairness with Adaptive Weights for Providers and Customers in Recommendation. *RecSys* (2024)
36. Yang, H., Liu, Z., Zhang, Z., Zhuang, C., Chen, X.: Towards robust fairness-aware recommendation. In: *RecSys* (2023)
37. Ye, X., Xu, C., Xu, J., Xie, X., Noah, H., Lab, A., Dong, Z., Wang, G.: Regret-aware Re-ranking for Guaranteeing Two-sided Fairness and Accuracy in Recommender Systems. *ACM Transactions on Information Systems* **1**(1) (4 2025)
38. Zehlike, M., Yang, K., Stoyanovich, J.: Fairness in ranking, part ii: Learning-to-rank and recommender systems. *ACM Computing Surveys* **55**(6), 1–41 (2022)