Multi-Attribute and Intersectional Exposure Fairness in Recommender Systems

Paul Anton Bachmann*
Technical University Vienna
Vienna, Austria
paul.bachmann@tuwien.ac.at

Hannah Eckert*
Johannes Kepler University Linz
Linz, Austria
Linz Institute of Technology
Linz, Austria
hannah.eckert@jku.at

Oleg Lesota
Johannes Kepler University Linz
Linz, Austria
oleg.lesota@jku.at

Martin Lackner
University of Applied Sciences
St. Pölten
St. Pölten, Austria
martin.lackner@ustp.at

Markus Schedl
Johannes Kepler University Linz
Linz Institute of Technology
Linz, Austria
markus.schedl@jku.at

Abstract

Ensuring fair exposure across demographic groups of content creators in recommender systems is crucial, as it affects the real-world visibility of minorities and underrepresented groups. Existing work on provider-side exposure fairness has primarily considered single demographic attributes, leaving multi-attribute and intersectional exposure fairness largely unexplored. Yet in real-world platforms, fairness concerns may arise simultaneously across multiple intersecting provider attributes. We introduce *multi-attribute* and *intersectional* perspectives on exposure fairness in recommender systems, propose an evaluation framework taking both perspectives into account, and present a post-processing approach for multi-attribute exposure fairness informed by online settings in social choice theory. Our experiments on real-world and synthetic datasets show a clear distinction between the two perspectives, as baseline approaches demonstrate a tendency to trade attribute-level fairness for improvements on the intersectional level. Our proposed method achieves a more favorable fairness-utility trade-off across attribute-wise exposure objectives and robust improvement of both exposure fairness aspects on the real world data. On synthetic data no method shows consistent improvement of both attribute-level and intersectional exposure fairness across all backbone models, highlighting the challenging nature of the task and the need to evaluate multi-attribute exposure fairness from both perspectives. With this work we aim to kick off the research on provider exposure fairness in recommender systems at a more granular level.

CCS Concepts

• **Information systems → Recommender systems**.

Keywords

Recommender Systems, Exposure Fairness, Position Bias

ACM Reference Format:
Paul Anton Bachmann, Hannah Eckert, Oleg Lesota, Martin Lackner, and Markus Schedl. 2026. Multi-Attribute and Intersectional Exposure Fairness in Recommender Systems. In *Proceedings of the 35th ACM International Conference on Information and Knowledge Management (CIKM '26), November 07–11, 2026, Rome, Italy*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3799682.3840023

1 Introduction

Visibility on online platforms has significant social consequences. It shapes the opportunities available to content creators and their experiences of marginalization [5]. Recommender systems can underrepresent content created by demographic groups that are already underrepresented in the catalog and historical interaction data [8, 14], thereby further reducing the visibility of marginalized groups. Exposure fairness is therefore concerned with the over- or underrepresentation of certain provider groups. An underexplored aspect in exposure fairness is that, in real world settings, unfairness can occur across several demographic attributes at once (e.g. an artist can face discrimination because of their race *and* their gender). Accumulating biases in machine learning models can also lead to intersectional effects, where disadvantages can intensify for those belonging to multiple disadvantaged groups [21]. Furthermore, a model can be intersectionally unfair, even if it appears fair with respect to each individual attribute. On the other hand, perfect intersectional fairness implies fairness across individual attributes. However, this goal may be unattainable. In the related setting of multidimensional political apportionment, proportional representation of intersectional groups of more than two attributes has even been shown to be impossible in general [2]. In addition, optimizing greedily toward intersectional fairness does not necessarily improve feature-wise fairness at each step and may even decrease it. We therefore argue that both intersectional and attribute-wise exposure fairness should be addressed and evaluated, which, to the best of our knowledge, has not yet been done in prior work.

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CIKM '26, Rome, Italy
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2539-5/2026/11
https://doi.org/10.1145/3799682.3840023

Our contributions are threefold. First, we introduce a joint perspective on intersectional and multi-attribute exposure fairness. Second, we adapt the disparate exposure metric [7] for the evaluation of multi-attribute and intersectional exposure fairness by scaling exposure deviations relative to group target proportions. Third, we propose a post-processing re-ranking method that explicitly mitigates attribute-level exposure imbalances, but prioritizes item providers that are underrepresented across multiple dimensions, thereby also accounting for intersectional effects.

2 Related Work

Although intersectional and multi-attribute exposure fairness has, to our knowledge, not yet been explicitly addressed in the context of recommender systems, there are multiple studies that address single-attribute exposure fairness in this context.

As user behavior is affected by position bias [3], i.e., the probability of an item being clicked depends on its position on the recommendation list, exposure fairness is commonly measured w.r.t. item ranking positions [4, 20]. Provider-side fairness has also been studied in settings that do not explicitly model position-dependent exposure [17], but we mostly focus on prior work that also accounts for ranking position. There are several post-processing approaches that regulate position-aware provider exposure by matching the distribution of a single provider attribute to a target distribution within each users recommendation list [12, 15, 16]. Such approaches can be adapted to intersectional exposure fairness by treating each intersectional group as a category of a single combined attribute. However, the number of such categories grows exponentially with the number of attributes, and enforcing the target distribution independently for each user limits the achievable exposure granularity. Gómez et al. [7] propose a post-processing approach that aims to achieve a global binary target exposure distribution, rather than enforcing the target distribution independently within each user's recommendation list. In [8], they extend their approach to non-binary attributes and in [9] they extend it to a general multi-objective recommendation setting, addressing both user- and provider perspectives. Wang et al. [22] and Wu et al. [23] also address multi-objective exposure fairness from provider- and user perspectives simultaneously. Multi-attribute provider exposure fairness can be viewed as a related multi-objective problem, as exposure must be balanced across several group definitions simultaneously. However, in our setting, these objectives arise within the provider side itself, between attribute-wise and intersectional group definitions. Because such objectives cannot always be satisfied within each individual top-k list, we view exposure fairness as an amortized allocation problem over a sequence of recommendation lists. This perspective connects our approach to work on online apportionment and sequential decision-making procedures for temporal fair allocations and perpetual voting in social choice theory [1, 2, 13].

3 Setting

We consider a setting where a set of users $\mathcal{U}$ each receive a ranked recommendation list of k items from a set of items $\mathcal{I}$. For each user u, a backbone model ϕ assigns a relevance score $R_\phi(u, i)$ to each item i, and outputs the top-k items ordered by relevance as recommendation list for each user. Our goal is to fairly allocate exposure to groups of providers across all users' recommendation lists. Provider groups are derived from item attributes. We consider multiple attributes simultaneously, resulting in intersectional groups whose exposure we later evaluate. Let X be an attribute with its corresponding values $\{X_1, X_2, \dots\}$ defining provider groups. The set of all attributes is denoted by $\mathcal{A}$. Each attribute partitions $\mathcal{I}$, i.e. for every $X \in \mathcal{A}$, each item belongs to exactly one group X_j of X. We write $X(i) = X_j$ to denote that item i belongs to the category X_j of attribute X. We define fairness as the degree to which provider group exposure complies with a target distribution. For each attribute considered (e.g., country or gender), a policy sets a target distribution of exposure for the different groups of that attribute. A policy may be externally given (e.g. by stakeholders) or derived from data. Denote by $f_X : \{X_1, X_2, \dots\} \to [0, 1]$ the density function for the target distribution of attribute X. It assigns a desired share of exposure to each group X_j. The number of provider groups who can receive exposure in a single recommendation list is limited by the number of recommendation ranks k. Representing any more than k groups fairly in each user's list is therefore impossible. Fairness notions focused on individual users only thus favor only larger provider groups, while smaller groups fail to be even considered. However, we can approximate the target exposure distributions globally *across all users*, similar to work in social choice theory on temporal fairness [1, 6, 13]. Following [16], we measure the exposure of an item ranked in position p as $1/\log_2(p+1)$. The exposure of a group X_j of attribute X in a recommendation list $\sigma = (\sigma_1, \dots, \sigma_k)$ of k items is denoted by $E_{X_j}(\sigma|k)$, the maximum exposure a category can receive in a list of k items by $E^*(k)$. The two are defined in Equation (1).

$$E_{X_j}(\sigma|k) = \sum_{p=1}^{k} \frac{\mathbb{1}(X(\sigma_p) = X_j)}{\log_2(p+1)} \quad E^*(k) = \sum_{p=1}^{k} \frac{1}{\log_2(p+1)} \tag{1}$$

4 Method: Exposure Deficit-Adjusted Relevance

To regulate the allocated exposure, we propose a post-processing method. We use exposure deficits of the different groups to inform adjustments of the user-item relevance scores produced by the backbone model. This is done sequentially, i.e. from user 1 to user $|\mathcal{U}|$. For each user, the relevance scores are adjusted according to the exposure distribution across all previous users. The allocated share of exposure up to a user u for a value X_j of an item attribute X, denoted as $\hat{f}_X^u$, is calculated as fraction of the total exposure. Here, $\sigma(u)$ denotes the recommendation list for a user u.

$$\hat{f}_X^u(X_j) = \frac{\sum_{u'=1}^{u} E_{X_j}(\sigma(u')|k)}{u \cdot E^*(k)} \tag{2}$$

As a measure of deficit (or surplus) in exposure for a group X_j up to user $u > 1$, we take the logarithm of the quotient of target and allocated exposure. We add a small $\varepsilon > 0$ to the allocated exposure to avoid division by zero for previously unrepresented groups.

$$\delta_{X_j}^u = \ln\left(\frac{f_X(X_j)}{\hat{f}_X^u(X_j) + \varepsilon}\right) \tag{3}$$

Intuitively, the exposure deficit $\delta_{X_j}^u$ is positive (resp. negative) if the group X_j has accumulated less (resp. more) than its target

exposure up to user u. The logarithm ensures equal penalization of over- and underrepresentation by the same factor. Based on the exposure deficits in the recommendations for all users $u' < u$, we adjust the relevance scores of the items for user u. For the first user, we use raw backbone relevance. To avoid recomputing all scores, we only consider a pre-selection of the top-100 most relevant items for each user. For each item, we aggregate the exposure deficits of all groups it belongs to. These are then used to compute the deficit-adjusted relevance score R^{δ}_{ϕ} with backbone model ϕ (cf. Equation (4)). The deficits are added up, weighted with a tradeoff parameter λ, and added to the base relevance score R_{ϕ}. This way, items belonging to multiple disadvantaged groups receive a larger score increase.

$$\mathrm{R}^{\delta}_{\phi}(u,i) = (1-\lambda)R_{\phi}(u,i) + \lambda \sum_{X \in \mathcal{A}} \delta^{u-1}_{X(i)} \tag{4}$$

The finalized recommendations list for the user is the ordered list of the top-k items w.r.t. deficit-adjusted relevance.

5 Experimental Setup

Evaluation Metrics. To evaluate both intersectional and multi-attribute exposure fairness, we adapt Disparate Exposure [7], by scaling it by group size. We define scaled disparate exposure for group X_j as the difference between overall allocated and target exposure, divided by target exposure:

$$SDE(X_j) = \frac{\hat{f}^{|\mathcal{U}|}_{X}(X_j) - f_X(X_j)}{f_X(X_j)} \tag{5}$$

Note that this can also be calculated for intersectional groups by adapting $\hat{f_X}$ and E_{X_j} to only measure the exposure of items that are in all respective groups of the intersecting attributes. We capture relative rather than absolute exposure differences, accounting for the fact that the same absolute exposure difference has a larger impact on smaller groups than on larger ones. To evaluate exposure fairness on a system level, we report the mean absolute scaled disparate exposure (MSDE) across all groups, considering attribute-wise or intersectional groups. We evaluate recommendation utility using normalized discounted cumulative gain (NDCG) at k=10.

Datasets. We evaluate our method using a real-world dataset with two and a synthetic dataset with three provider attributes. For the real world dataset, we use a subset of the LFM-2b dataset [19], where we use artist gender and country as attributes. We create the subset by retaining listening events from the year 2019, providers with both gender and country information, and countries with at least 100k listening events. The resulting dataset contains approximately 10k users, 1.5M items and 20M interactions. It includes three gender groups (female, male and non-binary) and 21 country groups, resulting in 63 intersectional groups. The demographic provider attributes are self reported. For the synthetic setting, we construct a dataset of 2M interactions generated for 20k users and 10k items. Each item has three attributes, each taking one of five possible values, resulting in 125 intersectional groups. User preferences are modeled using separate Dirichlet distributions over attributes. We sample separate preference distributions for a majority (90% of the users) and a minority (10% of the users), creating a setting in which majority preferences dominate minority preferences. The synthetic dataset is sampled with a fixed seed.

Backbone Models. As backbones for our post processing method as well as the baselines we use three common recommender system architectures: NeuMF [11], matrix factorization with Bayesian Personalized Ranking loss (BPR) [18] and LightGCN [10]. All methods are evaluated across all backbone models using three training seeds. Probabilistic post-processing methods are further evaluated with five post-processing seeds per training seed. We average reported results over all seeds. Exact training details and code, as well as the synthetic dataset are available at https://github.com/HannahEckert/MultiAttributeRecommenderSystems.

Target Distributions. We consider two different policies, interactions and population. For interactions, we aim for the proportions of the groups in the number of interactions. The goal of this policy is proportional representation of the chosen dataset. However, since datasets may contain bias and can structurally misrepresent real population proportions, we also consider the population policy for our real life dataset, where the target proportions are derived from the real-world composition of populations. For example, under the interactions policy, the target exposure share of female artists is approximately 0.3, corresponding to their share of interactions in the training data. A method allocating more than 30% of exposure to female artists would therefore count them as overrepresented relative to this policy. Under the population policy, the target exposure share for female artists is close to 0.5, so the same allocation may instead be considered underrepresentative. The interactions policy is often used in relevant literature [8, 14]. The population policy, on the other hand, aims at aligning the recommendation composition with real world proportions, avoiding any bias inherited from the dataset itself, resulting in a more drastic intervention.

Baselines. As baselines, we include the method introduced by Marras et al. [16], denoted Marras in the figures, and the method introduced by Kruse et al. [12], denoted NAILS in the figures, to represent approaches that match the distribution of a single provider attribute to a target distribution separately within each users recommendation list. To adapt these methods for intersectional fairness we treat each intersectional group as a category of a single combined attribute. We also include the method introduced by Gómez et al. [8], which we label Continent in the figures because the original work studies provider fairness across continents. We adapt it to the intersectional setting in the same way as Marras and NAILS. Unlike Marras and NAILS, it enforces the target distribution globally rather than separately for each user. The original Continent method does not explicitly specify a fairness-utility trade-off parameter. Since the method performs position swaps from a precomputed swap budget, we introduce a trade-off parameter that controls the fraction of swaps applied.

6 Results

We evaluate our method by comparing its trade-off between recommendation utility and exposure fairness across both intersectional and attribute-wise groups with single-attribute baselines adapted to the intersectional setting (Figures 1, 2 and 4). For the NeuMF and LightGCN backbones, we report results only for the interactions target distribution due to space constraints. Results for these

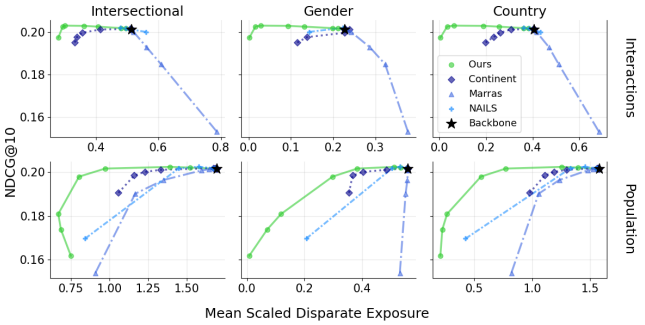


Figure 1: Fairness-utility trade-off on the LFM-2b dataset using the BPR backbone. We report NDCG@10 against MSDE for intersectional and attribute-wise provider groups under the interactions target distribution (top) and population target distribution (bottom). Points indicate different fairness weights, and black stars indicate the unmodified backbone model. Compared with the adapted single-attribute baselines, our method generally achieves lower MSDE across both group resolutions while remaining competitive in recommendation quality.

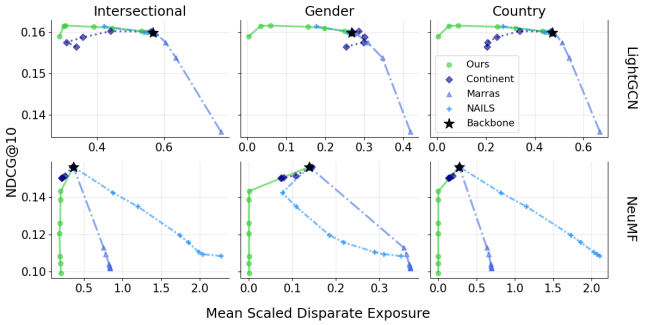


Figure 2: Fairness-utility trade-off on the LFM-2b dataset for the interactions target distribution using the LightGCN (top) and NeuMF (bottom) backbones. We report NDCG@10 against MSDE for intersectional and attribute-wise provider groups. Points correspond to different fairness weights, and black stars indicate the unmodified backbone models. Compared with the adapted single-attribute baselines, our method provides a more consistent reduction in MSDE across intersectional and attribute-wise groups.

backbone models under the population target distribution show similar trends and are available at https://github.com/HannahEckert/MultiAttributeRecommenderSystems.

For LFM-2b, our method generally reduces intersectional and attribute-wise MSDE while largely preserving recommendation performance (Figure 1 and 2). The performance drop is most pronounced under the population-based target distribution, which is expected because this target distribution differs more substantially from the test data distribution, on which the utility metrics is calculated. The Continent baseline also reduces MSDE across most group definitions. However our method achieves lower MSDE with comparable recommendation performance, as the final point reported for Continent corresponds to the maximum fairness weight supported by the method. Marras and NAILS perform less consistently in this setting. Under the interaction-based target distribution, they can even increase MSDE relative to the backbone. This suggests that enforcing fine grained fairness goals locally within individual top-k lists does not necessarily translate into lower MSDE across many provider groups. Figure 3, illustrates the effects of our method on individual intersectional and attribute-wise groups. Our method improves exposure fairness for all attribute-wise groups. It also reduces aggregate intersectional MSDE and improves exposure fairness for most intersectional groups. However, underrepresentation increases for a subset of intersectional groups. Similar group-level trade-offs are observed across the baselines, highlighting the importance of evaluating intersectional fairness beyond aggregate measures.

The synthetic setting (Figure 4) serves as a stress test for cases in which intersectional and attribute-wise fairness objectives are misaligned. In this setting, the adapted single-attribute baselines can achieve lower intersectional MSDE at a smaller utility cost in some cases. However, these reductions can come at the expense of substantially increased MSDE for attribute B. Our method is the only approach that does not substantially increase MSDE for attribute B across backbones. For LightGCN, however, the reduction in attribute-wise MSDE comes at the cost of higher intersectional MSDE for our method. No method consistently improves both

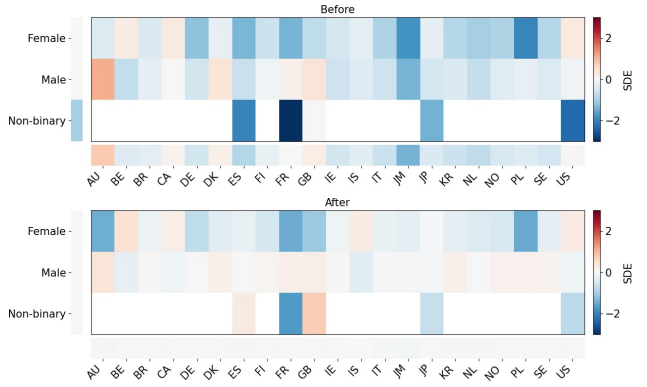


Figure 3: SDE for all intersectional groups of the LFM-2b dataset and the interactions target distribution using BPR backbone before (top) and after (bottom) applying our method with a fairness weight of $\lambda = 0.1$. We report SDE for all intersectional groups within the heatmap and attribute-wise SDE on the margins. Blue represents negative SDE (underrepresentation) and red represents positive SDE (overrepresentation). Our method decreased attribute-wise SDE to near zero for all investigated groups. However, while our method increases intersectional fairness at an aggregate level and for most intersectional groups, underrepresentation increases for some intersectional groups, such as Australian female creators.

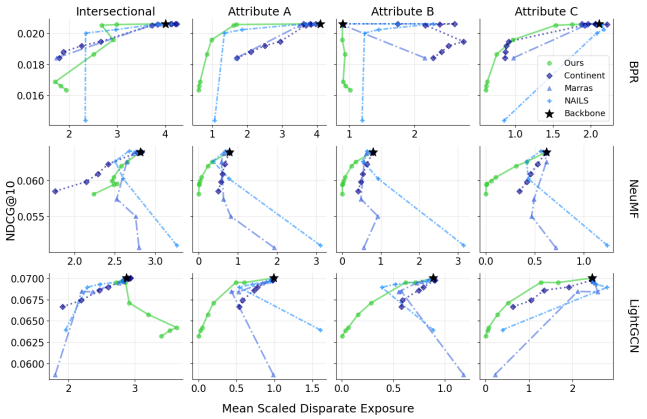


Figure 4: Fairness-utility trade-off on the synthetic dataset for the interactions target distribution. We report NDCG@10 against MSDE for intersectional and attribute-wise provider groups under the interactions target distribution. Points correspond to different fairness weights, and black stars indicate the unmodified backbone models. No method consistently improves both intersectional and attribute-wise MSDE across all backbone models.

intersectional and attribute-wise exposure fairness across all backbones. The preferred method therefore depends on the fairness priority. Under strong tension between objectives, our method may be preferable when attribute-wise exposure fairness should not be worsened, whereas the adapted Continent baseline may be preferable when intersectional exposure fairness is the primary objective. These results support our central claim that multi-attribute and intersectional exposure fairness need to be evaluated separately, as improvements under one perspective do not necessarily imply improvements under the other.

7 Conclusion and Future Work

This paper introduces a multi-attribute and intersectional perspective on provider-side exposure fairness in recommender systems. While prior approaches have primarily targeted fairness with respect to a single provider attribute, we argue that evaluating only individual attributes or even only intersectional groups can provide an incomplete picture on exposure fairness, as a method can improve one notion while worsening the other. To address this problem, we propose a post-processing framework inspired by online settings in social choice theory, which amortizes exposure across users and adjusts rankings according to accumulated exposure deficits across multiple provider attributes, favoring providers at the intersection of multiple underrepresented attributes. Our experiments on real-world and synthetic data show that intersectional and attribute-wise exposure fairness should not be treated as interchangeable objectives. On real world data, our method improves both intersectional and attribute-wise fairness at a more favorable fairness-utility trade-off than natural adaptations of single attribute baselines. On synthetic data, we observe strong tension between the fairness objectives and no method consistently improving attribute-level and intersectional exposure fairness across all backbone models. In this setting our method may be preferable when attribute-wise exposure fairness should not be worsened, whereas adapted baselines may be preferable when intersectional exposure fairness is the primary objective. Our method is therefore best understood as a mechanism for balancing exposure fairness across multiple demographic resolutions, favoring attribute-wise exposure fairness in case of strong tension between fairness objectives. Overall, our findings highlight the need to evaluate exposure fairness jointly across intersectional and attribute-wise group definitions.

One limitation of our work is that we do not consider user-side fairness concerns, which we plan to investigate in future work. Our method's sensitivity to changes in user ordering may play a role there, an aspect that we did not focus on for the goal of overall exposure fairness. Furthermore, while our method works well on the single real world dataset, in our synthetic data example none of the methods we compare consistently outperforms the others.

In future work, we plan to further investigate the tension between attribute-wise and intersectional exposure fairness and extend the framework to additional perspectives, such as user-side or individual artist fairness. Furthermore, we want to study the setting from a theoretical perspective, investigating the (im)possibility of perfect intersectional fairness.

Acknowledgments

This research was supported by the State of Upper Austria, and was partly funded by the Linz Institute of Technology (LIT). In addition, the research was funded in part by the Austrian Science Fund (FWF): Cluster of Excellence Bilateral Artificial Intelligence (https://doi.org/10.55776/COE12) and the doc.funds.connect project Human-Centered Artificial Intelligence (https://doi.org/10.55776/DFH23). We are grateful to the reviewers for pointing out numerous possible improvements for the final version of this paper.

References

[1] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. 2018. How to Make Envy Vanish Over Time. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, Ithaca NY USA, 593–610. doi:10.1145/3219166.3219179

[2] Javier Cembrano, Jose Correa, Svenja M. Griesbach, and Victor Verdugo. 2026. Online Proportional Apportionment. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 4846–4860. doi:10.1137/1.9781611978971.174

[3] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*. 87–94.

[4] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 275–284.

[5] Brooke Erin Duffy and Colten Meisner. 2023. Platform governance at the margins: Social media creators' experiences with algorithmic (in) visibility. *Media, Culture & Society* 45, 2 (2023), 285–304.

[6] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Temporal Fairness in Multiwinner Voting. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 20 (March 2024), 22633–22640. doi:10.1609/aaai.v38i20.30273

[7] Elizabeth Gómez, Ludovico Boratto, and Maria Salamó. 2021. Disparate impact in item recommendation: A case of geographic imbalance. In *European Conference on Information Retrieval*. Springer, 190–206.

[8] Elizabeth Gómez, Ludovico Boratto, and Maria Salamó. 2022. Provider fairness across continents in collaborative recommender systems. *Information Processing & Management* 59, 1 (2022), 102719.

[9] Elizabeth Gómez, David Contreras, Ludovico Boratto, and Maria Salamó. 2024. MOREGIN: multi-objective recommendation at the global and individual levels.

In *European Conference on Information Retrieval*. Springer, 21–38.

[10] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.

[11] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.

[12] Johannes Kruse, Kasper Lindskow, Michael Riis Andersen, Ryotaro Shimizu, Julian McAuley, Pierre-Alexandre Mattei, and Jes Frellsen. 2025. Normative Alignment of Recommender Systems via Internal Label Shift. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*. 1240–1245.

[13] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 02 (April 2020), 2103–2110. doi:10.1609/aaai.v34i02.5584

[14] Oleg Lesota, Emilia Parada-Cabaleiro, Stefan Brandl, Elisabeth Lex, Navid Rekabsaz, and Markus Schedl. 2022. Traces of Globalization in Online Music Consumption Patterns and Results of Recommendation Algorithms.. In *ISMIR*. 291–297.

[15] Ramon Lopes, Rodrigo Alves, Antoine Ledent, Rodrygo LT Santos, and Marius Kloft. 2024. Recommendations with minimum exposure guarantees: A post-processing framework. *Expert Systems with Applications* 236 (2024), 121164.

[16] Mirko Marras, Ludovico Boratto, Guilherme Ramos, and Gianni Fenu. 2022. Regulating group exposure for item providers in recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on research and development in information retrieval*. 1839–1843.

[17] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the web conference 2020*. 1194–1204.

[18] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. 452–461.

[19] Markus Schedl, Stefan Brandl, Oleg Lesota, Emilia Parada-Cabaleiro, David Penz, and Navid Rekabsaz. 2022. LFM-2b: A Dataset of Enriched Music Listening Events for Recommender Systems Research and Fairness Analysis. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) *(CHIIR '22)*. Association for Computing Machinery, 337–341.

[20] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*. Association for Computing Machinery, New York, NY, USA, 2219–2228. doi:10.1145/3219819.3220088

[21] Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. *Advances in neural information processing systems* 32 (2019).

[22] Chenyang Wang, Yankai Liu, Yuanqing Yu, Weizhi Ma, Min Zhang, Yiqun Liu, Haitao Zeng, Junlan Feng, and Chao Deng. 2023. Two-sided calibration for quality-aware responsible recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 223–233.

[23] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on research and development in information retrieval*. 703–714.