



HAL
open science

How Fair is Your Diffusion Recommender Model?

Daniele Malitesta, Giacomo Medda, Erasmo Purificato, Ludovico Boratto,
Fragkiskos D. Malliaros, Mirko Marras, Ernesto William de Luca

► **To cite this version:**

Daniele Malitesta, Giacomo Medda, Erasmo Purificato, Ludovico Boratto, Fragkiskos D. Malliaros, et al.. How Fair is Your Diffusion Recommender Model?. 2024. hal-04849064

HAL Id: hal-04849064

<https://inria.hal.science/hal-04849064v1>

Preprint submitted on 19 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

How Fair is Your Diffusion Recommender Model?

Daniele Malitesta¹, Giacomo Medda², Erasmo Purificato³,
Ludovico Boratto², Fragkiskos D. Malliaros¹, Mirko Marras²,
Ernesto William De Luca^{4,5}

¹Université Paris-Saclay, CentraleSupélec, Inria, Centre for Visual Computing (CVN), Gif-sur-Yvette, France.

²University of Cagliari, Cagliari, Italy.

³Joint Research Centre, European Commission, Ispra, Italy.

⁴Otto von Guericke University Magdeburg, Magdeburg, Germany.

⁵Leibniz Institute for Educational Media, Georg Eckert Institute, Brunswick, Germany.

Contributing authors: daniele.malitesta@centralesupelec.fr;
giacomo.medda@unica.it; erasmo.purificato@ec.europa.eu;
ludovico.boratto@acm.org; fragkiskos.malliaros@centralesupelec.fr;
mirko.marras@acm.org; ernesto.deluca@ovgu.de;

Abstract

Diffusion-based recommender systems have recently proven to outperform traditional generative recommendation approaches, such as variational autoencoders and generative adversarial networks. Nevertheless, the machine learning literature has raised several concerns regarding the possibility that diffusion models, while learning the distribution of data samples, may inadvertently carry information bias and lead to unfair outcomes. In light of this aspect, and considering the relevance that fairness has held in recommendations over the last few decades, we conduct one of the first fairness investigations in the literature on *DiffRec*, a pioneer approach in diffusion-based recommendation. First, we propose an experimental setting involving DiffRec (and its variant L-DiffRec) along with nine state-of-the-art recommendation models, two popular recommendation datasets from the fairness-aware literature, and six metrics accounting for accuracy and consumer/provider fairness. Then, we perform a twofold analysis, one assessing models' performance under accuracy and recommendation fairness *separately*, and the other identifying if and to what extent such metrics can strike a performance *trade-off*. Experimental results from both studies confirm the initial

unfairness warnings but pave the way for how to address them in future research directions.

Keywords: Diffusion Models, Fairness, Recommender Systems

1 Introduction and motivation

Research in recommender systems (RSs) has recently been shaped by the latest advances in deep learning, *e.g.*, graph neural networks [1], large language models [2], and diffusion models [3]. While the first two families have already received significant attention in the literature, diffusion-based RSs are still in early stages.

As of today, diffusion models [4–6] have settled as the leading solution in generative AI [7] and revolutionized the paradigm of generating new (but realistic) data with several applications in computer vision [4, 8], natural language processing [9, 10], graph data [11, 12], chemistry, and bioinformatics [13, 14]. The core principle behind diffusion models lies in their twofold learning procedure, involving a **forward** process where samples from a real-data distribution are iteratively corrupted by adding noise, *e.g.*, Gaussian noise, and a **reverse** process where the model is trained to reconstruct the original data by denoising the corrupted samples.

As observed by Wang et al. [3], the diffusion model framework is well-suited for recommendation settings, where it may effectively learn to denoise users’ implicit feedback containing noisy interactions, *i.e.*, false positives. Their approach, called *DiffRec*, shows improved performance to traditional generative RSs, *e.g.*, those based on variational autoencoders [15] and generative adversarial networks [16] by addressing their limited representational power (the former) or unstable training procedures (the latter). Similarly, other works have proposed to leverage diffusion models for specific recommendation tasks such as point-of-interests [17], sequential [18, 19], or knowledge graph-based recommendation [20].

Diffusion models have achieved widespread success in many domains and applications, but they are a relatively recent trend in machine learning, still facing several open challenges. For instance, some works have raised concerns regarding **unfairness** that could affect diffusion models. As observed in generative AI, **bias** is intrinsic in real-world data, and diffusion models might generate **unfair** outcomes while learning from such a **biased** data. In their seminal paper, Friedrich et al. [21] discuss how stable diffusion methods are prone to generate synthetic images that lack diversity, *e.g.*, firefighters are only associated with white males. Similarly, other works, such as Bianchi et al. [22], discuss how textual prompts with trait annotations may produce stereotypical images, while Cho et al. [23] observe bias in the generated images in terms of people color tones and occupations. Recent papers have proposed possible countermeasures, such as introducing fair guidance instructions during the data generation process [21], obfuscating the sensitive attributes [24], or adopting an ad-hoc fine-tuning [25].

In light of the above, a natural question arises: “*Can diffusion models be harmful to recommendation fairness likewise?*”. The question is pertinent given the large relevance that **bias** and **fairness** have held in recommendations for decades [26–30].

Thus, we consider it is imperative to rigorously evaluate the performance of popular diffusion-based RSs, with a specific focus on potential **unfairness issues**, as there is currently **no literature on this topic**. Indeed, it is worth underlining that such recommendation models are still in their pioneer era. A thorough investigation of their performance can help identify and address observed pitfalls.

To this end, in the current work, we select one of the pioneer approaches in diffusion-based recommendation, DiffRec [3] (and its variant L-DiffRec), to ultimately answer the following research question: “*Is DiffRec a fair recommender system compared to the state-of-the-art?*”. To conduct our investigation, we consider two datasets (with users’ attribute data) and nine additional state-of-the-art recommender systems. Our purpose is to benchmark their performance on six recommendation metrics accounting for **accuracy**, **consumer** [31–33], and **provider fairness** [27, 34, 35]. Concretely, we run a twofold performance analysis, where we first measure each performance metric **separately**, and then in a **trade-off** setting. We believe this twofold perspective can help identify performance patterns by encompassing all actors (consumers and providers) involved in the recommendation scenario [32].

Results depict a worrying scenario, outlining that **fairness concerns raised in machine learning could also hold for the recommendation domain**. Nonetheless, introducing additional components in the diffusion-based process, as for L-DiffRec, seems to mitigate the negative effects of unfairness; indeed, these findings pave the way to novel future directions in diffusion-based recommendation where fairness is preserved. Despite being a **preliminary investigation**, to the best of our knowledge, this is the first attempt in the literature to provide a rigorous and fairness-aware evaluation perspective on diffusion-based recommendation.

2 Background

This section provides the background notions regarding the two main topics of this paper, namely, diffusion-based recommendation (with a special focus on DiffRec [3]) and the evaluation metrics to assess fairness (consumer and provider sides) in recommendation.

2.1 Diffusion recommender model

Similarly to diffusion models, DiffRec [3] proposes to treat recommendation as a denoising process of historical user-item interactions, which may be extremely noisy, *i.e.*, contain false positives. In detail, the recommendation model is trained to first corrupt users’ interactions through noise (**forward** process) and then reconstruct them via an ad-hoc denoising procedure (**reverse** process). Differently from traditional diffusion models employed on images, users’ historical interactions are not corrupted into pure noise, but the ratio of added noise decreases consistently.

Let \mathcal{U} and \mathcal{I} be the sets of users and items in the system, respectively, with $\mathbf{x}_u \in \{0, 1\}^{|\mathcal{I}|}$ as the vector of user-item interactions for user u . Each vector entry $\mathbf{x}_{u,i} = 1$ denotes user u interacted with item i , otherwise $\mathbf{x}_{u,i} = 0$. Then, the forward process, in a Markov chain process, corrupts the user interaction vector into T steps by gradually adding Gaussian noise. The formulation for any step $0 \leq t \leq T$ is: $q(\mathbf{x}_t | \mathbf{x}_{t-1}) =$

$\mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$, where (with a notation abuse) we indicate and set $\mathbf{x}_0 = \mathbf{x}_u$, $\mathbf{x}_0 \in \mathbb{R}$, at the initial step, while \mathcal{N} stands for the Gaussian distribution, $\beta_t \in (0, 1)$ controls the magnitude of the Gaussian noise added at each step t , and \mathbf{I} is the standard deviation of the distribution. Conversely, the reverse process involves training any machine learning model, *e.g.*, a neural network, to approximate the mean and standard deviation of the data distribution and recover the original user interaction vector. Again, for any step t , we have: $p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t))$, where $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$ and $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)$ are the learned mean and standard deviation through the machine learning model having parameters θ .

Furthermore, to address large-scale recommendation (a pretty common scenario in real-world applications), Wang et al. [3] also propose L-DiffRec, a lighter-weight version of DiffRec. This variant first clusters pre-trained item embeddings, *e.g.*, from LightGCN, into C clusters. Afterward, the initial user interaction vector \mathbf{x}_0 is split into C sub-vectors $\{\mathbf{x}_0^c\}_{c=1}^C$ according to the corresponding c -th item cluster. Then, a variational autoencoder is trained to compress each of user interaction sub-vectors \mathbf{x}_0^c , so that the diffusion, *i.e.*, forward/reverse processes, occurs in a low-dimensional latent space, ideally in parallel computation across all splits. Finally, the original user interaction vector is reconstructed through the decoding process of the previous variational autoencoder.

Note that in Wang et al. [3], the authors eventually propose a second variant to DiffRec, called T-DiffRec, which can encode the temporal information during the training. However, T-DiffRec is specialized for sequential recommendation, which would require a dedicated study to compare its performance with other analogous models. To this end, we exclude T-DiffRec from our experiments, leaving its adoption to future extensions of our study.

Our analysis aims to offer a **preliminary** outline on the fairness performance of diffusion-based recommender systems. Although other works have recently contributed with novel diffusion-based solutions [17–20], DiffRec could be rightly considered as the **pioneer** approach in this domain for chronological reasons. In light of this, we selected only this model and plan to extend our investigation to other diffusion models in forthcoming works.

2.2 Fairness evaluation in recommendation

Recent studies focused on assessing fairness under the notion of *demographic parity* [31, 33, 35]. On the one hand [35, 36], we operationalize demographic parity on the *consumer* side as the absolute difference in recommendation utility between users’ groups. We select the respective fairness counterparts of nDCG and Recall, namely Δ nDCG and Δ Recall. On the other hand, *provider* fairness is estimated as the absolute difference in exposure between items’ groups (Δ Exp) [34, 35] and the average percentage of long-tail items in the top- k lists (APLT) [27, 32]. Lower values of these metrics reflect fairer recommendations, except for APLT. For page constraints, we do not dive deep into their formulations. The interested reader may refer to the above-cited papers for a detailed presentation.

3 Experimental study

In this section, we present our proposed experimental study aimed at assessing fairness in diffusion-based recommendation. First, we describe the baselines and datasets adopted for the analysis. Then, we provide reproducibility details for the conducted experiments.

3.1 Baselines

We select nine state-of-the-art recommendation models along with DiffRec and L-DiffRec: BPRMF [37], ItemkNN [38], NeuMF [39], LightGCN [40], UltraGCN [41], XSimGCL [42], EASE [43], MultiVAE [15], and RecVAE [44]. As reported in Table 1, our selection depends on the following criteria: (i) the baselines are well-established solutions, accepted at top-tier conferences and journals; (ii) they cover a diverse set of recommendation families, such as graph neural networks and generative models; (iii) many approaches have been already analyzed under recommendation fairness, highlighting an existing literature gap in diffusion-based recommendation.

3.2 Datasets

We assess the performance of the selected baselines on MovieLens-1M (**ML1M**) [45] and Foursquare Tokyo (**FTKY**) [46], which are frequently adopted in fairness-aware literature in recommendation [33, 36, 47]. While ML1M is a popular recommendation dataset that collects movies and users’ preferences on them, FTKY involves interactions represented by users’ check-ins at points of interest in Tokyo. In both datasets, users’ sensitive information comes in a binary gender label. Thus, fairness analyses on the consumer side adopt the gender attribute included in both datasets, whereas provider fairness is assessed on a binary partition of the items’ set, namely *short-head* (popular) and *long-tail* (niche) items. Following [27, 34, 35], we select the 20% most popular items as short-head and the rest as long-tail. Dataset statistics are reported in Table 2.

3.3 Reproducibility

Datasets are split into train, validation, and test sets (70%/10%/20%) following the temporal hold-out method. Then, for the models training and evaluation, we rely on the framework RecBole [48] for training and evaluating all the baselines and, for the two diffusion-based recommender systems, we adopt their original code implementations¹. All recommendation baselines are trained with a grid-search exploration of the hyper-parameters by following the original papers and selecting the hyper-parameter configuration providing the highest value of the Recall@20 on the validation set.

4 Results and discussion

In this section, we aim to answer the following research question: “*Is DiffRec a fair recommender system compared to the state-of-the-art?*”. To this end, we conduct a

¹<https://github.com/YiyangXu/DiffRec>.

Table 1 Model baselines indicating publication year, venue, and fairness-aware studies that analyzed them.

Models	Year	Venue	Analyzed fairness in
BPRMF [37]	2009	UAI	[26–29, 31, 32]
ItemkNN [38]	2001	WWW	[26, 30, 31, 36]
NeuMF [39]	2017	WWW	[26–28, 30, 31, 36]
LightGCN [40]	2020	SIGIR	[29, 32, 36]
UltraGCN [41]	2021	CIKM	[32]
XSimGCL [42]	2024	TKDE	x
EASE [43]	2019	WWW	[26]
MultiVAE [15]	2018	WWW	[26, 49]
RecVAE [44]	2020	WSDM	x
DiffRec [3]	2023	SIGIR	x
L-DiffRec [3]	2023	SIGIR	x

Table 2 Recommendation datasets along with their statistics and information about users’ groups and attributes.

Datasets	ML1M [45]	FTKY [46]
# Users	6,040	7,240
# Items	3,650	5,779
# Interactions	1,000,130	353,769
Sparsity	95.46%	99.15%
Gender Repr.	M : 71.7%; F : 28.3%	M : 87.9%; F : 12.1%
Min. Interactions	20	20

twofold analysis. First, we investigate the accuracy and fairness performance *separately*, and then we measure the *trade-off* among performance metrics to examine whether DiffRec is suitable to provide both fair and useful recommendations.

4.1 Single-metric performance

Table 3 shows the results of the selected baselines on ML1M and FTKY in terms of recommendation accuracy, *i.e.*, Recall and nDCG, consumer fairness, *i.e.*, Δ Recall and Δ nDCG, and provider fairness, *i.e.*, APLT and Δ Exp. We recall that all metrics are calculated on top-20 recommendation lists, and are expressed as percentages (%).

On the recommendation *accuracy* level, the results are almost aligned with the findings from the related literature, where graph- and diffusion-based recommender systems generally dominate over the other baselines on both recommendation datasets. DiffRec shows slightly superior performance than its variant L-DiffRec, but the latter is more computationally lightweight. This gap in performance between the two DiffRec variants might be due to the diffusion process being conducted in the latent space by L-DiffRec. Specifically, the variational autoencoders employed by L-DiffRec compress the users’ preferences signals, which might lead to information loss and a non-perfect reconstruction in the decoding process. On the other hand, DiffRec forward and reverse

Table 3 Results on top-20 recommendation lists for the selected baselines for ML1M and FTKY when measuring recommendation accuracy, *i.e.*, Recall and nDCG, consumer fairness, *i.e.*, Δ Recall and Δ nDCG, and provider fairness, *i.e.*, APLT and Δ Exp. For each metric, the arrow direction indicates whether higher (\uparrow) or lower (\downarrow) values stand for better performance. Boldface and underline refer to best and second-to-best results. All metric values are expressed as percentages (%).

	Metrics	BPRMF	ItemkNN	NeuMF	LightGCN	UltraGCN	XSimGCL	EASE	MultiVAE	RecVAE	DiffRec	L-DiffRec
ML1M	Recall (\uparrow)	10.67	9.87	10.44	10.16	10.74	10.46	6.45	8.25	9.68	<u>10.71</u>	10.38
	nDCG (\uparrow)	12.94	12.22	12.67	<u>13.09</u>	12.82	12.80	8.62	12.04	12.86	13.19	12.85
	Δ Recall (\downarrow)	0.43	1.08	0.32	0.19	0.06	<u>0.14</u>	0.70	0.48	0.53	0.86	0.66
	Δ nDCG (\downarrow)	1.57	1.98	1.54	1.43	<u>1.37</u>	1.34	1.66	2.43	2.07	2.03	2.25
	APLT (\uparrow)	<u>12.47</u>	8.01	11.75	8.11	13.73	10.50	9.00	2.63	5.38	6.18	10.81
	Δ Exp (\downarrow)	<u>86.29</u>	91.18	87.36	91.32	84.98	88.78	90.54	97.02	94.10	93.67	87.26
FTKY	Recall (\uparrow)	12.17	11.77	12.48	11.03	12.22	<u>13.07</u>	6.19	10.06	12.30	13.10	12.48
	nDCG (\uparrow)	10.73	10.48	10.93	9.87	10.85	<u>11.37</u>	5.48	9.21	10.90	11.53	10.77
	Δ Recall (\downarrow)	0.94	0.69	<u>0.59</u>	0.77	1.05	0.98	0.40	0.97	1.00	0.80	0.61
	Δ nDCG (\downarrow)	0.22	0.19	<u>0.08</u>	0.47	0.29	0.23	0.10	0.59	0.29	0.22	0.02
	APLT (\uparrow)	9.93	18.27	10.20	7.42	5.36	11.06	6.32	3.72	7.27	10.75	<u>16.97</u>
	Δ Exp (\downarrow)	89.27	78.91	89.39	91.87	94.32	88.25	93.56	95.75	92.03	88.57	<u>80.85</u>

processes fully capture the discrete preference information, without dealing with the potential loss of details caused by information bottleneck.

Moving to the experiments on *fairness*, results exhibit distinct trends across the datasets. On ML1M, graph-based approaches such as UltraGCN and XSimGCL are able to provide high-utility and fair recommendations, while the diffusion models settle among the unfairest approaches across all the baselines. On FTKY, the fairest recommender systems are instead the generative ones, such as EASE and L-DiffRec, and the traditional ones, such as ItemKNN and NeuMF. The diverse trends in consumer fairness across ML1M and FTKY might derive from their different domains. Specifically, users’ interactions in the movie domain are driven by intrinsically individual preferences towards specific movie categories, such as those sharing the same actors or the same genre, whereas venue check-ins in FTKY might derive from goings out decided as groups, *e.g.*, with friends or partners. Biases in movies’ preferences might be exacerbated by the RSs compared with venues’ preferences.

The *consumer unfairness* exhibited by DiffRec on ML1M underscores that diffusion models are prone to increase the bias impact, as also emphasized by other studies in different domains. However, a common observation across the datasets is that the L-DiffRec variant reports fairer recommendations than DiffRec and the other family-related generative models, *e.g.*, MultiVAE and RecVAE. Particularly, regarding APLT and Δ Exp, L-DiffRec exhibits fairer outcomes on the *provider* side. This might be attributed to the intrinsic clustering operation of L-DiffRec, which shifts the learning process from following the common popularity bias to capturing more nuanced patterns across items’ categories.

4.2 Trade-off analysis

As the three analyzed properties, *i.e.*, accuracy, consumer, and provider fairness, represent different and possibly *contrasting* sides of the recommendation scenario, we extend our previous analysis by assessing the trade-off among them, as recently done by related studies [30, 32]. Indeed, we believe the current analysis may suitably complement the findings from the *single-metric* study conducted above, motivating and explaining from other complementary perspectives the previous performance trends.

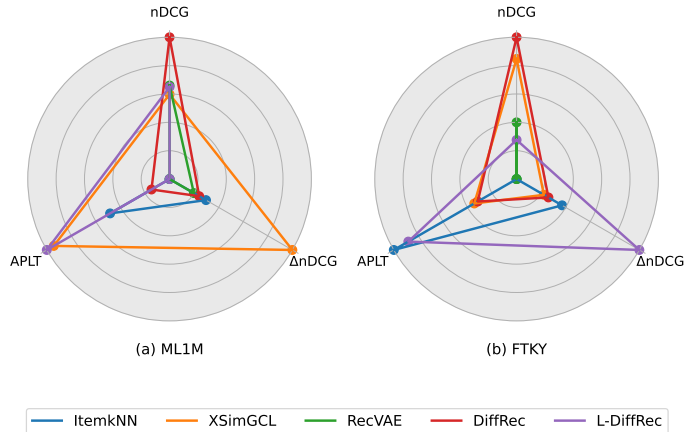


Fig. 1 Kiviati diagrams of the trade-off performance for five selected baselines on recommendation accuracy (nDCG), consumer fairness (Δ nDCG), and provider fairness (APLT). For better visualization, metric values have been normalized in $[0, 1]$, and Δ nDCG has been further rescaled to adhere to the principle “the higher the better”.

The Kiviati diagrams in Figure 1 represent the performance for both datasets of five representative baselines, *i.e.*, ItemkNN, XSimGCL, RecVAE, DiffRec, and L-DiffRec, on three metrics, *i.e.*, nDCG, Δ nDCG, and APLT, accounting for accuracy, consumer, and provider fairness. To provide an easier interpretation of the plots, we follow [32] and normalize the metrics into $[0, 1]$ through min-max normalization and rescale Δ nDCG to adhere to the principle that “the higher the values the better the performance”.

On the ML1M dataset, most of the findings from the *single-metric* analysis are generally confirmed. Indeed, we observe that XSimGCL is the most balanced baseline across all models, being capable of providing an optimal trade-off among all metrics. Conversely, ItemkNN, RecVAE, and DiffRec cannot generally increase one metric side without harming the other ones. Interestingly, though, while ItemkNN tends to maximize fairness over the accuracy, RecVAE and DiffRec, *i.e.*, generative and diffusion-based models, do the opposite. This becomes especially evident on DiffRec, where the model is largely improving accuracy at the detriment of provider/consumer fairness. The exception to the trend is L-DiffRec, reaching the second-best trade-off soon after XSimGCL. In fact, on FTKY, L-DiffRec exhibits even the best trade-off compared to the selected baselines.

Thus, the findings from the *trade-off* analysis, with a more clarifying visualization, confirm the unfairness warnings raised in the *single-metric* study but, at the same time, suggest there might be room for improvement in diffusion-based recommendation toward more fairness-aware recommendations.

5 Conclusion and future work

This work settles as a preliminary attempt to rigorously assess the performance of diffusion-based recommender systems, specifically focusing on one of the pioneer

approaches, DiffRec, and its variant L-DiffRec. The relevance of our study is emphasized by the machine learning literature that has recently observed that the intrinsic bias within the data may cause diffusion models to generate unfair outcomes. To this end, we decided to tailor our investigation to the performance of (L-)DiffRec under common fairness properties in recommender systems. Specifically, we conducted an extensive analysis involving two recommendation datasets (with sensitive information), nine state-of-the-art recommendation systems, and six recommendation metrics accounting for accuracy, consumer, and provider fairness. First, in a single-metrics evaluation, results generally outlined that DiffRec might exhibit unfairness issues; however, such issues are mitigated by the L-DiffRec variant, especially on the provider side. Then, a complementary trade-off analysis among accuracy and fairness properties confirmed that L-DiffRec might be the key to reaching optimal metrics balancing. We believe the take-home message is that diffusion models cannot be plugged into recommender systems as they are, since they would bring (and even exacerbate) the unfairness issues already observed in machine learning. However, with careful and tailored modifications to the rationale and models' architecture, *e.g.*, L-DiffRec, the issues might be properly addressed. To confirm our final assumption, future works will aim to extend our study to the other DiffRec variant, T-DiffRec, as well as other diffusion-based recommender systems in the literature.

References

- [1] Wu, S., Sun, F., Zhang, W., Xie, X., Cui, B.: Graph neural networks in recommender systems: A survey. *ACM Comput. Surv.* **55**(5), 97–19737 (2023)
- [2] Wu, L., Zheng, Z., Qiu, Z., Wang, H., Gu, H., Shen, T., Qin, C., Zhu, C., Zhu, H., Liu, Q., Xiong, H., Chen, E.: A survey on large language models for recommendation. *CoRR* **abs/2305.19860** (2023)
- [3] Wang, W., Xu, Y., Feng, F., Lin, X., He, X., Chua, T.: Diffusion recommender model. In: *SIGIR*, pp. 832–841. ACM, ??? (2023)
- [4] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *NeurIPS* (2020)
- [5] Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., Yang, M.: Diffusion models: A comprehensive survey of methods and applications. *ACM Comput. Surv.* **56**(4), 105–110539 (2024)
- [6] Cao, H., Tan, C., Gao, Z., Xu, Y., Chen, G., Heng, P.-A., Li, S.Z.: A survey on generative diffusion models. *IEEE Transactions on Knowledge and Data Engineering*, 1–20 (2024)
- [7] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial networks. *Commun. ACM* **63**(11), 139–144 (2020)

- [8] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR, pp. 10674–10685. IEEE, ??? (2022)
- [9] Austin, J., Johnson, D.D., Ho, J., Tarlow, D., Berg, R.: Structured denoising diffusion models in discrete state-spaces. In: NeurIPS, pp. 17981–17993 (2021)
- [10] Li, X., Thickstun, J., Gulrajani, I., Liang, P., Hashimoto, T.B.: Diffusion-lm improves controllable text generation. In: NeurIPS (2022)
- [11] Niu, C., Song, Y., Song, J., Zhao, S., Grover, A., Ermon, S.: Permutation invariant graph generation via score-based generative modeling. In: AISTATS. Proceedings of Machine Learning Research, vol. 108, pp. 4474–4484. PMLR, ??? (2020)
- [12] Liu, C., Fan, W., Liu, Y., Li, J., Li, H., Liu, H., Tang, J., Li, Q.: Generative diffusion models on graphs: Methods and applications. In: IJCAI, pp. 6702–6711. ijcai.org, ??? (2023)
- [13] Xu, M., Yu, L., Song, Y., Shi, C., Ermon, S., Tang, J.: Geodiff: A geometric diffusion model for molecular conformation generation. In: ICLR. OpenReview.net, ??? (2022)
- [14] Hoogeboom, E., Satorras, V.G., Vignac, C., Welling, M.: Equivariant diffusion for molecule generation in 3d. In: ICML. Proceedings of Machine Learning Research, vol. 162, pp. 8867–8887. PMLR, ??? (2022)
- [15] Liang, D., Krishnan, R.G., Hoffman, M.D., Jebara, T.: Variational autoencoders for collaborative filtering. In: WWW, pp. 689–698. ACM, ??? (2018)
- [16] Wang, J., Yu, L., Zhang, W., Gong, Y., Xu, Y., Wang, B., Zhang, P., Zhang, D.: IRGAN: A minimax game for unifying generative and discriminative information retrieval models. In: SIGIR, pp. 515–524. ACM, ??? (2017)
- [17] Qin, Y., Wu, H., Ju, W., Luo, X., Zhang, M.: A diffusion model for POI recommendation. *ACM Trans. Inf. Syst.* **42**(2), 54–15427 (2024)
- [18] Li, Z., Sun, A., Li, C.: Diffurec: A diffusion model for sequential recommendation. *CoRR* **abs/2304.00686** (2023)
- [19] Ma, H., Xie, R., Meng, L., Chen, X., Zhang, X., Lin, L., Kang, Z.: Plug-in diffusion model for sequential recommendation. In: AAAI, pp. 8886–8894. AAAI Press, ??? (2024)
- [20] Jiang, Y., Yang, Y., Xia, L., Huang, C.: Diffkg: Knowledge graph diffusion model for recommendation. In: WSDM, pp. 313–321. ACM, ??? (2024)
- [21] Friedrich, F., Schramowski, P., Brack, M., Struppek, L., Hintersdorf, D., Luccioni, S., Kersting, K.: Fair diffusion: Instructing text-to-image generation models on

fairness. CoRR **abs/2302.10893** (2023)

- [22] Bianchi, F., Kalluri, P., Durmus, E., Ladhak, F., Cheng, M., Nozza, D., Hashimoto, T., Jurafsky, D., Zou, J., Caliskan, A.: Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In: FAccT, pp. 1493–1504. ACM, ??? (2023)
- [23] Cho, J., Zala, A., Bansal, M.: DALL-EVAL: probing the reasoning skills and social biases of text-to-image generation models. In: ICCV, pp. 3020–3031. IEEE, ??? (2023)
- [24] Choi, Y., Park, J., Kim, H., Lee, J., Park, S.: Fair sampling in diffusion models through switching mechanism. In: AAI, pp. 21995–22003. AAAI Press, ??? (2024)
- [25] Shen, X., Du, C., Pang, T., Lin, M., Wong, Y., Kankanhalli, M.S.: Finetuning text-to-image diffusion models for fairness. CoRR **abs/2311.07604** (2023)
- [26] Anelli, V.W., Bellogín, A., Noia, T.D., Jannach, D., Pomo, C.: Top-n recommendation algorithms: A quest for the state-of-the-art. In: UMAP, pp. 121–131. ACM, ??? (2022)
- [27] Mansoury, M., Abdollahpouri, H., Pechenizkiy, M., Mobasher, B., Burke, R.: A graph-based approach for mitigating multi-sided exposure bias in recommender systems. *ACM Trans. Inf. Syst.* **40**(2), 32–13231 (2022)
- [28] Sun, J., Guo, W., Zhang, D., Zhang, Y., Regol, F., Hu, Y., Guo, H., Tang, R., Yuan, H., He, X., Coates, M.: A framework for recommending accurate and diverse items using bayesian graph convolutional neural networks. In: KDD, pp. 2030–2039. ACM, ??? (2020)
- [29] Zhao, M., Wu, L., Liang, Y., Chen, L., Zhang, J., Deng, Q., Wang, K., Shen, X., Lv, T., Wu, R.: Investigating accuracy-novelty performance for graph-based collaborative filtering. In: SIGIR, pp. 50–59. ACM, ??? (2022)
- [30] Boratto, L., Fenu, G., Marras, M., Medda, G.: Practical perspectives of consumer fairness in recommendation. *Inf. Process. Manag.* **60**(2), 103208 (2023)
- [31] Boratto, L., Fenu, G., Marras, M., Medda, G.: Consumer fairness in recommender systems: Contextualizing definitions and mitigations. In: ECIR (1). *Lecture Notes in Computer Science*, vol. 13185, pp. 552–566. Springer, ??? (2022)
- [32] Anelli, V.W., Deldjoo, Y., Noia, T.D., Malitesta, D., Paparella, V., Pomo, C.: Auditing consumer- and producer-fairness in graph collaborative filtering. In: ECIR (1). *Lecture Notes in Computer Science*, vol. 13980, pp. 33–48. Springer, ??? (2023)

- [33] Vassøy, B., Langseth, H.: Consumer-side fairness in recommender systems: a systematic survey of methods and evaluation. *Artif. Intell. Rev.* **57**(4), 101 (2024)
- [34] Boratto, L., Fabbri, F., Fenu, G., Marras, M., Medda, G.: Robustness in fairness against edge-level perturbations in gnn-based recommendation. In: *ECIR* (3). *Lecture Notes in Computer Science*, vol. 14610, pp. 38–55. Springer, ??? (2024)
- [35] Naghiaei, M., Rahmani, H.A., Deldjoo, Y.: Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In: *SIGIR*, pp. 770–779. ACM, ??? (2022)
- [36] Boratto, L., Fabbri, F., Fenu, G., Marras, M., Medda, G.: Counterfactual graph augmentation for consumer unfairness mitigation in recommender systems. In: *CIKM*, pp. 3753–3757. ACM, ??? (2023)
- [37] Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: BPR: bayesian personalized ranking from implicit feedback. In: *UAI*, pp. 452–461. AUAI Press, ??? (2009)
- [38] Sarwar, B.M., Karypis, G., Konstan, J.A., Riedl, J.: Item-based collaborative filtering recommendation algorithms. In: *WWW*, pp. 285–295. ACM, ??? (2001)
- [39] He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.: Neural collaborative filtering. In: *WWW*, pp. 173–182. ACM, ??? (2017)
- [40] He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M.: Lightgcn: Simplifying and powering graph convolution network for recommendation. In: *SIGIR*, pp. 639–648. ACM, ??? (2020)
- [41] Mao, K., Zhu, J., Xiao, X., Lu, B., Wang, Z., He, X.: Ultragen: Ultra simplification of graph convolutional networks for recommendation. In: *CIKM*, pp. 1253–1262. ACM, ??? (2021)
- [42] Yu, J., Xia, X., Chen, T., Cui, L., Hung, N.Q.V., Yin, H.: Xsimgel: Towards extremely simple graph contrastive learning for recommendation. *IEEE Trans. Knowl. Data Eng.* **36**(2), 913–926 (2024)
- [43] Steck, H.: Embarrassingly shallow autoencoders for sparse data. In: *WWW*, pp. 3251–3257. ACM, ??? (2019)
- [44] Shenbin, I., Alekseev, A., Tutubalina, E., Malykh, V., Nikolenko, S.I.: Recvae: A new variational autoencoder for top-n recommendations with implicit feedback. In: *WSDM*, pp. 528–536. ACM, ??? (2020)
- [45] Harper, F.M., Konstan, J.A.: The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.* **5**(4), 19–11919 (2016)
- [46] Yang, D., Zhang, D., Qu, B.: Participatory cultural mapping based on collective

behavior data in location-based social networks. *ACM Trans. Intell. Syst. Technol.* **7**(3), 30–13023 (2016)

- [47] Liu, Q., Mu, L., Sugumaran, V., Wang, C., Han, D.: Pair-wise ranking based preference learning for points-of-interest recommendation. *Knowl. Based Syst.* **225**, 107069 (2021)
- [48] Zhao, W.X., Hou, Y., Pan, X., Yang, C., Zhang, Z., Lin, Z., Zhang, J., Bian, S., Tang, J., Sun, W., Chen, Y., Xu, L., Zhang, G., Tian, Z., Tian, C., Mu, S., Fan, X., Chen, X., Wen, J.: Recbole 2.0: Towards a more up-to-date recommendation library. In: *CIKM*, pp. 4722–4726. ACM, ??? (2022)
- [49] Togashi, R., Abe, K., Saito, Y.: Scalable and provably fair exposure control for large-scale recommender systems. In: *WWW*, pp. 3307–3318. ACM, ??? (2024)