

Debiasing Neighbor Aggregation for Graph Neural Network in Recommender Systems

Minseok Kim
Amazon Alexa AI
Seattle, WA, USA
kimseok@amazon.com

Jinoh Oh
Amazon Alexa AI
Seattle, WA, USA
ojino@amazon.com

Jaeyoung Do
Amazon Alexa AI
Seattle, WA, USA
domjae@amazon.com

Sungjin Lee
Amazon Alexa AI
Seattle, WA, USA
sungjinl@amazon.com

ABSTRACT

Graph neural networks (GNNs) have achieved remarkable success in recommender systems by representing users and items based on their historical interactions. However, little attention was paid to GNN’s vulnerability to *exposure bias*: users are exposed to a limited number of items so that a system only learns a biased view of user preference to result in suboptimal recommendation quality. Although inverse propensity weighting is known to recognize and alleviate exposure bias, it usually works on the final objective with the model outputs, whereas GNN can also be biased during neighbor aggregation. In this paper, we propose a simple but effective approach, neighbor aggregation via inverse propensity (NAVIP) for GNNs. Specifically, given a user-item bipartite graph, we first derive propensity score of each user-item interaction in the graph. Then, inverse of the propensity score with Laplacian normalization is applied to debias neighbor aggregation from exposure bias. We validate the effectiveness of our approach through our extensive experiments on two public and Amazon Alexa datasets where the performance enhances up to 14.2%.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender systems, graph neural networks, exposure bias, collaborative filtering, inverse propensity scoring

ACM Reference Format:

Minseok Kim, Jinoh Oh, Jaeyoung Do, and Sungjin Lee. 2022. Debiasing Neighbor Aggregation for Graph Neural Network in Recommender Systems. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22)*, October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3511808.3557576>

1 INTRODUCTION

Graph neural networks (GNNs) became a standard approach in collaborative filtering [6, 12, 16] by aggregating information of interacted users and items (i.e., neighbors) in a user-item bipartite graph to obtain meaningful performance gain [5, 22, 24]. However,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '22, October 17–22, 2022, Atlanta, GA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9236-5/22/10...\$15.00
<https://doi.org/10.1145/3511808.3557576>

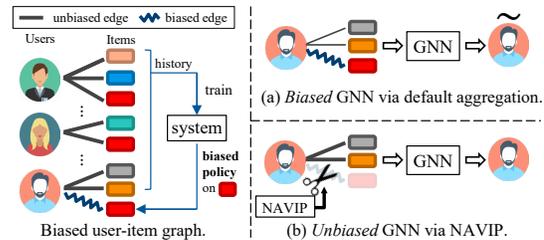


Figure 1: Comparison of (a) *biased default GNN* and (b) *debiased GNN via NAVIP* given a *biased user-item graph*.

as shown in Figure 1(a), real-world user-item graphs are often biased because edges are user responses to the system behavior decided by a policy. Specifically, a system can expose only some part of all items to users so that how user-item interactions occur becomes biased by the system’s policy. Since GNNs represent a node by aggregating neighbors connected via edges, GNNs induce a biased view of user preference, which may *not* be aligned with true user interest. As such, recommender systems based on GNNs leave room for improvement through debiasing the neighbor aggregation process with respect to the exposure bias.

A common way to mitigate exposure bias includes inverse propensity scoring (IPS) that emphasizes less-exposed items during training [3, 4, 9, 17, 19–21]. Specifically, IPS multiplies the inverse propensity of user-item interactions in the objective function, which accordingly upweights the learning on less-exposed items. However, since IPS deals with the bias problem *after* GNNs encode users and items into embeddings, the system is still susceptible to the bias caused by the unadjusted process of neighbor aggregation.

In this paper, we propose a simple yet effective approach, neighbor aggregation via inverse propensity (NAVIP) for GNNs. Specifically, given a user-item bipartite graph, we first derive the propensity score of each user-item interaction in the graph. We then use the inverse propensity score with Laplacian normalization as edge weight for the neighbor aggregation process. As such, less popular neighbors are relatively emphasized in an embedding to balance out the biased local structure of each target node as shown in Figure 1(b). We validate the effectiveness of NAVIP with a comprehensive experiment on three real-world datasets including a large-scale voice assistant system, Amazon Alexa.

2 PRELIMINARY AND RELATED WORK

In online services, user-item interactions data $\{(u, i, Y_{u,i})\} \in \mathcal{D}_{\pi_0}$ is collected according to *policy* π_0 , which is a prior on user behaviors. *Recommender systems* perform personalized recommendation by finding similar users and items based on the collected data \mathcal{D}_{π_0} , which is known as collaborative filtering [6, 12, 16]. Specifically, given a recommender model $\phi(\cdot)$ and its prediction for user u ’s

relevance on item i as $\phi(h_u, h_i)$ based on their embeddings, h_u and h_i , we minimize empirical risk $R(\pi_0)$ on implicit feedback $Y_{u,i}$,

$$R(\pi_0) = \sum_{(u,i,Y_{u,i}) \in \mathcal{D}_{\pi_0}} \mathcal{L}(\phi(h_u, h_i), Y_{u,i}), \quad (1)$$

where \mathcal{L} is a loss function such as BPR [15], and $Y_{u,i}$ is 1 if the interaction between user u and item i is observed or 0 if not.

A prevalent approach for collaborative filtering follows two steps: (i) transforms each user and item into a latent representation, h_u and h_i , then (ii) exploits the representations to estimate the relevance $\phi(h_u, h_i)$ based on dot product [15] or non-linear neural networks (DNN) [6] based on h_u and h_i .

2.1 GNN for Recommender System

GNN is the state-of-the-art approach for learning representation over graphs [5]. The biggest difference between GNN and conventional CF methods is that GNN explicitly leverages neighbors of a target node during inference of the target node's representation. Specifically, given a bipartite graph of users and items as nodes, the embedding h of a user u at layer l is computed as follows,

$$h_u^l = \text{TRANSFORM}(h_u^{l-1}, a_u^l) \quad (2)$$

$$a_u^l = \text{AGGREGATE}(h_i^{l-1}, i \in \mathcal{N}(u)), \quad (3)$$

where a_u^l denotes the aggregated embedding of u 's neighbors $\mathcal{N}(u)$ at layer l . Note that we can augment the embedding of item i in the same way. The design of $\text{AGGREGATE}(\cdot)$ and $\text{TRANSFORM}(\cdot)$, or leveraging high-order connectivity are used to distinguish GNN methods. Initial GNN methods [2, 10, 13] used only one-hop neighbors using one GNN layer, while following studies [22, 24] exploit multiple layers with non-linear DNN to benefit from high-order connectivity. Recently, LIGHTGCN [5] claimed that such non-linear DNN is not essential for GNN, and showed a state-of-the-art performance by employing symmetric Laplacian norm in $\text{AGGREGATE}(\cdot)$ and identity function for $\text{TRANSFORM}(\cdot)$,

$$h_u^l = a_u^l = \sum_{i \in \mathcal{N}(u)} \frac{h_i^{l-1}}{\sqrt{|\mathcal{N}(u)| |\mathcal{N}(i)|}}. \quad (4)$$

2.2 Inverse Propensity Scoring

In a service, users are exposed to only a part of items selected by a system policy π_0 . Thus, collected log data \mathcal{D}_{π_0} becomes dependent toward the policy π_0 , which is called *exposure bias*. Naturally, a recommender ϕ trained on \mathcal{D}_{π_0} becomes biased and even intensifies the bias [23, 25]. To this end, IPS is popularly exploited to acquire unbiased estimator $R(\pi)$ of the performance on the uniform item exposure policy π by reweighting \mathcal{D}_{π_0} with propensity [17, 18],

$$\hat{R}(\pi; \mathcal{D}_{\pi_0}) = \frac{1}{|\mathcal{D}_{\pi_0}|} \sum_{(u,i,Y_{u,i}) \in \mathcal{D}_{\pi_0}} \frac{\mathcal{L}(\phi(h_u, h_i), Y_{u,i})}{p(Y_{u,i}=1|\pi_0)}, \quad (5)$$

where $p(Y_{u,i}=1|\pi_0)$ is propensity defined in Definition 2.1.

Definition 2.1. Propensity $p(Y_{u,i}=1|\pi_0)$ is the observation probability of user-item interaction $Y_{u,i}$ under policy π_0 .

However, IPS has a few drawbacks, which are complemented by several studies. For example, SNIPS [20] was suggested to address high variance of IPS at the expense of a small estimation bias [8, 19].

Recently, UEBPR [3] improves IPS-based recommender learning by reflecting recommendation explainability during training. However, we note that the above-described IPS studies work in the objective function *after* GNN neighbor aggregation, therefore the system is still biased even with IPS.

3 METHODOLOGY

Previous methods have mainly focused on applying IPS on loss as shown in Eq. (5), but gave less attention to how the neighbors are aggregated even though it is one of the key factors that contribute to the success of GNN. Because collected log data \mathcal{D}_{π_0} has exposure bias (i.e., $\mathcal{D}_{\pi_0} \neq \mathcal{D}_{\pi}$), GNNs also employ neighbor aggregation to biased interactions $\{x = (u, i, Y_{u,i}) \mid x \in \mathcal{D}_{\pi_0} \wedge x \notin \mathcal{D}_{\pi}\}$, which can bias user embedding h_u^l . Even worse, the bias can propagate to users without biased interactions due to the repeated GNN layers.

3.1 NAVIP for Debiasing GNN

Inspired by the motivation, we suggest NAVIP that applies IPS weighting scheme to neighbor aggregation to reduce the impact of biased interactions and debias GNN neighbor aggregation. That is, NAVIP can be considered as upweighting interactions with low propensity and downweighting interactions with high propensity during aggregation. The rationale behind NAVIP is that neighbors with lower propensity interaction reveal more about the true interest of a given target user because such interactions are not likely to occur by a biased system policy. On the other hand, because interactions with high propensity often occur as a response to a biased system policy, they may not be the true user interest.

For simplicity, we tentatively assume a simple neighbor aggregation, $a_u^l = \mathbb{E}[h_i^{l-1}]$, which is the expected embedding of neighbors at layer $l-1$. For clarity, we further define $a_u^l(\pi)$ to be the aggregated neighbor embedding of user u under the system policy π , and $a_u^l(\pi|\mathcal{D}_{\pi})$ to be the *Monte-Carlo approximated* aggregated neighbor embedding of π over \mathcal{D}_{π} . Formally speaking,

$$a_u^l(\pi) \approx \frac{1}{|\mathcal{N}_{\pi}(u)|} \sum_{i \in \mathcal{N}_{\pi}(u)} h_i^{l-1}(\pi) = a_u^l(\pi|\mathcal{D}_{\pi}) \quad (6)$$

where $h_i^{l-1}(\pi)$ is a node embedding at $l-1$ layer under policy π . In practice, $a_u^l(\pi)$ needs to be approximated based on \mathcal{D}_{π_0} because \mathcal{D}_{π} is not available while developing π . NAVIP estimates $a_u^l(\pi)$ with respect to \mathcal{D}_{π} by combining Eq. (5) and Eq. (6),

$$a_u^l(\pi) \approx \frac{1}{|\mathcal{N}_{\pi_0}(u)|} \sum_{i \in \mathcal{N}_{\pi_0}(u)} \frac{h_i^{l-1}(\pi)}{p(Y_{u,i}=1|\pi_0)} = a_u^l(\pi|\mathcal{D}_{\pi_0}). \quad (7)$$

However, we note that using inverse propensity results in numerical instability issue due to unnormalized magnitude of neighbor weights. To avoid the issue, NAVIP further normalizes neighbor aggregation by replacing $\frac{1}{|\mathcal{N}_{\pi_0}(u)|}$,

$$\frac{1}{Z_{\pi}(u)} \sum_{i \in \mathcal{N}_{\pi_0}(u)} \frac{h_i^{l-1}(\pi)}{p(Y_{u,i}=1|\pi_0)} = a_u^l(\pi|\mathcal{D}_{\pi_0}), \quad (8)$$

where $Z_{\pi}(u)$ is a normalizing term defined by $\sum_{i \in \mathcal{N}_{\pi_0}(u)} \frac{1}{p(Y_{u,i}=1|\pi_0)}$. We replace Eq. (3) with Eq. (8) to debias neighbor aggregation.

3.1.1 *Connection to Laplacian.* The neighbor aggregation function of NAVIP can be formulated in a matrix form. Specifically,

$$\Pi_{u,i} = \begin{cases} \frac{1}{P(Y_{u,i}=1|\pi_0)} & \text{if } Y_{u,i} = 1 \\ 0 & \text{else,} \end{cases} \quad (9)$$

$$\Lambda = \begin{bmatrix} \mathbf{0} & \mathbf{\Pi} \\ \mathbf{\Pi}^\top & \mathbf{0} \end{bmatrix}$$

Then, the weighted random walk Laplacian can be shown as

$$\mathbf{A}^l(\pi_0|\mathcal{D}_\pi) = \mathbf{D}^{-1}\mathbf{\Lambda}\mathbf{H}^l(\pi_0), \quad (10)$$

where \mathbf{A}^l and \mathbf{H}^l are the matrix of aggregated neighbor embedding and node embedding at layer l , $\mathbf{\Lambda}$ is the adjacency matrix, and \mathbf{D} is a diagonal degree matrix whose k -th diagonal elements is the sum of k -th row of $\mathbf{\Lambda}$, $D_{kk} = \sum_i \Lambda_{ki}$.

It is worth noting that $\mathbf{D}^{-1}\mathbf{\Lambda}$ can be interpreted as random walk Laplacian, L^w , on a *weighted graph*, whose weights are inverse propensity weights. This simple difference from default aggregation encourages neighbors with low propensity edge become dominant in h_u^l , even with more GNN layers.

3.2 Model Training

The simplicity of NAVIP brings two benefits during training: (i) *no computational overhead* due to NAVIP during training because NAVIP debiases GNN neighbor aggregation *without* any additional training parameter, and (ii) *no objective adjustment* because NAVIP modifies neighbor aggregation before final model output $\phi(h_u, h_i)$ so as to be independent of learning objectives as in Eq.(1) or Eq.(5). In this regard, NAVIP can follow given recommender’s training scheme while debiasing the recommender. In addition, it is worth noting that NAVIP can achieve the debiasing effect for the pretrained model *without* additional training (see Section 4.2 for more details).

4 EVALUATION

We conducted experiments to answer the following questions:

- Does NAVIP improve overall recommender performance?
- How does NAVIP impact recommender’s exposure bias?
- Does NAVIP have overlapping effect with IPS?

4.1 Experimental Settings

4.1.1 *Datasets.* We conducted experiments on three real-world datasets: Alexa, Yahoo¹, and Movielens².

- **Alexa**: it contains a truncated user interaction logs from a voice assistant service, Amazon Alexa. The interactions are collected for a week in August 2021. Redundant user-item interactions such as users’ daily weather checking were filtered out to align setting with other two datasets.
- **Yahoo**: it contains users’ five-star song ratings from Yahoo! Music. Different from other two datasets, Yahoo contains unbiased test data which is obtained by asking 5, 400 users to rate 10 randomly selected songs. To make the ratings as implicit feedback, we transformed all the ratings into $Y_{u,i} = 1$ and unobserved entries as $Y_{u,i} = 0$ to indicate whether a user u interacted with an item i .
- **Movielens**: it contains users’ five-star ratings for movies. Like Yahoo, all the ratings were binarized.

¹<http://webscope.sandbox.yahoo.com/>

²<http://grouplens.org/datasets/movielens/>

Dataset	Users	Items	Interactions
Alexa	40k	2.5k	0.7m
Yahoo	15k	1.0k	0.3m
Movielens	943	1.6k	0.1m

Table 1: Summary of three real-world datasets.

We filtered out users and items that contain less than 10 interactions following usual practice [6]. Table 1 shows the profile of the three datasets after preprocessing.

We created pseudo-unbiased testing and validation data for Alexa and Movielens dataset. Because real-world datasets are biased, test data created by random splitting is not effective to evaluate the debiasing effect. Having a dedicated unbiased test dataset generated from pure random exposure is ideal, but this is practically infeasible in many cases. Following [3, 17], we generated pseudo-unbiased test data by sampling 5% of total data based on the inverse of relative item popularity. In case of Yahoo, we used the original test data, but for validation set, we followed the same process to generate a unbiased validation data.

4.1.2 *Implementation detail.* We used LIGHTGCN [5]³, which is the state-of-the-art GNN-based recommender system, as a backbone GNN throughout the experiment. We investigated three neighbor aggregations to validate the effect of neighbor aggregation.

- **MEAN**: a basic neighbor aggregation strategy that averages embeddings of neighbors.
- **PROPENSITY**: a bias-oriented neighbor aggregation strategy with propensity $P(Y_{u,i} = 1|\pi_0)$ as edge weight.
- **NAVIP**: a *debiasing* neighbor aggregation strategy that uses IPS as edge weight for each neighbor as in Eq.(8).

In addition, to validate whether the effect of NAVIP overlaps with that of IPS methods, we conducted experiment with the state-of-the-art IPS method for ranking loss, UEBPR [3].

Following [3, 17], we estimate the propensity of an item to a user by relative item popularity,

$$p(Y_{u,i} = 1|\pi_0) = \sqrt{\frac{\sum_{u \in U} Y_{u,i}}{\max_{l \in I} \sum_{u \in U} Y_{u,l}}}, \quad (11)$$

where U and I are the set of users and items, respectively.

All the experiments were trained for 100 epochs with batch size 256. The size of input feature and embedding are 64. We used Adam [11] with a learning rate $\eta = 0.003$. Our implementation was written using Pytorch and tested on Nvidia Tesla V100. We conducted experiment 10 times and reported their average.

4.1.3 *Evaluation Metrics.* We used two widely-recognized metrics, hit rate (HR) and normalized discounted cumulative gain (NDCG). Given a model’s recommendation list, HR measures the proportion of true user-item interactions in a recommended list, and NDCG additionally reflects the ranking of interactions. Two metrics were calculated with top- k recommendation: a model chooses k most probable items that each user is expected to like and suggest them to the user. We varied the number of recommended items $k \in \{5, 10, 20\}$. Because it is time-consuming to rank all items, a recommender evaluated the rank of the user’s 1 interacted item among randomly sampled 99 non-interacted items following literature [7].

³<https://github.com/gusye1234/LightGCN-PyTorch>

Dataset	Method	HR			NDCG		
		@5	@10	@20	@5	@10	@20
Alexa	MEAN	0.266	0.364	0.498	0.190	0.222	0.256
	PROPENSITY	0.153	0.255	0.410	0.096	0.129	0.167
	NAVIP	0.287	0.379	0.509	0.217	0.247	0.279
Yahoo	MEAN	0.145	0.242	0.397	0.093	0.124	0.163
	PROPENSITY	0.134	0.224	0.376	0.086	0.115	0.153
	NAVIP	0.148	0.249	0.408	0.095	0.127	0.167
Movielens	MEAN	0.350	0.506	0.681	0.233	0.283	0.327
	PROPENSITY	0.343	0.493	0.668	0.230	0.278	0.322
	NAVIP	0.349	0.511	0.691	0.230	0.283	0.328

Table 2: Recommendation performance of pretrained LIGHTGCN varying neighbor aggregation method during testing.

Dataset	Method	HR			NDCG		
		@5	@10	@20	@5	@10	@20
Alexa	MEAN	0.266	0.364	0.498	0.190	0.222	0.256
	PROPENSITY	0.260	0.357	0.489	0.187	0.218	0.251
	NAVIP	0.268	0.368	0.504	0.192	0.224	0.258
Yahoo	MEAN	0.145	0.242	0.397	0.093	0.124	0.163
	PROPENSITY	0.145	0.241	0.392	0.093	0.124	0.162
	NAVIP	0.147	0.246	0.397	0.095	0.126	0.164
Movielens	MEAN	0.350	0.506	0.681	0.233	0.283	0.327
	PROPENSITY	0.347	0.503	0.678	0.233	0.283	0.327
	NAVIP	0.352	0.512	0.687	0.234	0.286	0.330

Table 3: Recommendation performance of each neighbor aggregation method during training and testing.

4.2 Overall Performance

Table 2 shows the performance of pretrained LIGHTGCN model with varying neighbor aggregation during inference to evaluate the effect of neighbor aggregation. The performance difference among neighbor aggregation methods verifies that neighbor aggregation *significantly* affects the quality of embedding so as to impact on the success of recommendation. Overall, NAVIP shows the best performance among three neighbor aggregations in most cases.

Specifically, NAVIP outperforms MEAN and PROPENSITY by up to 14.2% and 126% in terms of $NDCG@5$ in Alexa dataset, respectively. Such performance improvement is attributed to the debiasing effect of NAVIP on neighbor aggregation, which successfully downweights neighbors with biased interaction so as to receive relatively more information from relevant neighbors. This claim can be further supported by the performance of PROPENSITY which always recorded the worst performance among three neighbor aggregations. That is, PROPENSITY highlights neighbors with biased interactions, which may not represent the true interest of users.

We also investigate the impact of neighbor aggregation when they are used in both training and inference (Table 3). The performance gap among neighbor aggregation strategies decreased so that all three aggregations showed performance similar to MEAN. We conjecture that training learnable LIGHTGCN input features under the same learning objective can weaken the effect of different neighbor aggregation. Nevertheless, the order of recommendation performance among neighbor aggregation strategies maintains; NAVIP scored the best performance in general, while PROPENSITY consistently performs the worst.

Data	Method	Head Items			Tail Items		
		HR@5	HR@10	HR@20	HR@5	HR@10	HR@20
Alexa	Mean	0.268	0.369	0.505	0.264	0.361	0.494
	NAVIP	0.271	0.373	0.510	0.265	0.363	0.499
Yahoo	Mean	0.497	0.638	0.795	0.033	0.071	0.141
	NAVIP	0.483	0.631	0.781	0.038	0.076	0.159
Movielens	Mean	0.649	0.812	0.933	0.077	0.142	0.236
	NAVIP	0.647	0.811	0.931	0.080	0.149	0.250

Table 4: Accuracy on head (top 10% popularity) and tail (bottom 10% popularity) items with respect to two neighbor aggregations, MEAN and NAVIP.

Dataset	Method	HR			NDCG		
		@5	@10	@20	@5	@10	@20
Alexa	MEAN	0.299	0.400	0.533	0.218	0.251	0.284
	NAVIP	0.301	0.403	0.538	0.220	0.253	0.287
Yahoo	MEAN	0.152	0.250	0.408	0.098	0.130	0.169
	NAVIP	0.153	0.252	0.411	0.098	0.130	0.170
Movielens	MEAN	0.350	0.512	0.699	0.230	0.282	0.330
	NAVIP	0.349	0.512	0.693	0.229	0.282	0.328

Table 5: Recommendation performance with UEBPR on two neighbor aggregation methods, MEAN and NAVIP. The results that exceeds the performance in Table 3 are marked in bold.

4.3 Impact on Recommender’s Exposure Bias

Table 4 shows the performance of LIGHTGCN on head items with top 10% popularity and tail items with bottom 10% popularity. The result shows that the performance on head items is higher than on tail items for all datasets, because recommenders are known to expose popular items more to users [1, 14]. NAVIP alleviates such issues by emphasizing user interest in tail items, which leads to more exposure to tail items. That is, NAVIP encourages LIGHTGCN to aggregate more information from neighbors with a low propensity, which correspond to tail items, so that LIGHTGCN can reveal user interest in tail items. Although NAVIP may lead to a performance drop on head items as in Yahoo and Movielens dataset, it is negligible because NAVIP improves overall performance.

4.4 Relationship with IPS

Table 5 shows the performance of LIGHTGCN trained with IPS method UEBPR during training. While NAVIP with IPS showed the best performance in Alexa and Yahoo datasets for both types of Laplacian norm, NAVIP without IPS scored the best in Movielens dataset. We point out that such inconsistent results in Movielens arose from the high variance of IPS that may lead training to be far from the ideal training objective $R(\pi)$ [20]. Nevertheless, we can conclude that NAVIP and IPS can be synergetic in debiasing GNN.

5 CONCLUSION

In this paper, we proposed NAVIP to debias in exposure bias during neighbor aggregation of GNN. NAVIP acquires inverse propensity score for each user-item interaction and exploits its Laplacian norm as neighbor weight. Experiments conducted on three real-world datasets confirmed that NAVIP can successfully debias GNN neighbor aggregation and enhance performance up to 14.2%.

REFERENCES

- [1] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286* (2019).
- [2] Rianne van den Berg, Thomas N Kipf, and Max Welling. 2018. Graph convolutional matrix completion. *KDD Workshop on Deep Learning Day* (2018).
- [3] Khalil Damak, Sami Khenissi, and Olfa Nasraoui. 2021. Debaised Explainable Pairwise Ranking from Implicit Feedback. In *RecSys*. 321–331.
- [4] Alois Gruson, Praveen Chandar, Christophe Charbuillet, James McInerney, Samantha Hansen, Damien Tardieu, and Ben Carterette. 2019. Offline evaluation to make decisions about playlist recommendation algorithms. In *ICWSM*. 420–428.
- [5] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [6] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *TheWebConf*. 173–182.
- [7] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *SIGIR*. 549–558.
- [8] Tim Hesterberg. 1995. Weighted average importance sampling and defensive mixture distributions. *Technometrics* 37, 2 (1995), 185–194.
- [9] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased learning-to-rank with biased feedback. In *ICWSM*. 781–789.
- [10] Santosh Kabbur, Xia Ning, and George Karypis. 2013. Fism: factored item similarity models for top-n recommender systems. In *SIGKDD*. 659–667.
- [11] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [12] Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. 1997. Grouplens: Applying collaborative filtering to usenet news. *Commun. ACM* 40, 3 (1997), 77–87.
- [13] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *KDD*. 426–434.
- [14] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback loop and bias amplification in recommender systems. In *CIKM*. 2145–2148.
- [15] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. *UAI* (2009).
- [16] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. 1994. Grouplens: An open architecture for collaborative filtering of netnews. In *CSCW*. 175–186.
- [17] Yuta Saito. 2019. Unbiased Pairwise Learning from Implicit Feedback. In *NeurIPS 2019 Workshop on Causal Machine Learning*.
- [18] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. 2020. Unbiased recommender learning from missing-not-at-random implicit feedback. In *WSDM*. 501–509.
- [19] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*. PMLR, 1670–1679.
- [20] Adith Swaminathan and Thorsten Joachims. 2015. The self-normalized estimator for counterfactual learning. *NeurIPS*.
- [21] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to rank with selection bias in personal search. In *SIGIR*. 115–124.
- [22] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*. 165–174.
- [23] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. 2018. Unbiased offline recommender evaluation for missing-not-at-random implicit feedback. In *RecSys*. 279–287.
- [24] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *KDD*. 974–983.
- [25] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. *SIGIR* (2021).