Diversified and Accurate Recommendations for a Series of Users

Jongjin Kim
Seoul National University
Seoul, Korea
j2kim99@snu.ac.kr

U Kang
Seoul National University
Seoul, Korea
ukang@snu.ac.kr

ABSTRACT

When we sequentially recommend top-k items to users, how can we recommend them diversely while maintaining accuracy? Aggregate-level diversity is an important topic to maximize the potential profit of platforms by exposing a variety of items. However, existing aggregately diversified recommender systems ignore the order of users and assume that all recommendations happen at once. In reality, users receive recommendations sequentially so latter users' information is not given during serving former users.

In this work, we propsoe the problem of sequentially diversified recommendation and SAPID to address the problem. SAPID removes the popularity bias from the model through a negative sampling mechanism based on temporal popularities. Then, SAPID decides which items to recommend immediately or later according to their estimated exposure opportunities. Extensive experiments show that SAPID shows the state-of-the-art performance in real-world datasets by achieving up to 4.25 times higher accuracy and up to 3.92% increased diversity compared to the best competitor.

CCS CONCEPTS

• **Information systems** → **Recommender systems**.

KEYWORDS

diversified recommendation, sequential recommendation, aggregate-level diversity

ACM Reference Format:

Jongjin Kim and U Kang. 2024. Diversified and Accurate Recommendations for a Series of Users . In *Proceedings of 30TH ACM SIGKDD CONFERENCE ON KNOWLEDGE DISCOVERY AND DATA MINING (SIGKDD'24).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

When we sequentially recommend top-k items to users, how can we increase the diversity among the items while maintaining the accuracy? Recommender systems become an essential element for users to discover items to consume [5, 7, 11]. Hence, it is important to expose as many products as possible to improve the total profit of the platform [2, 4]. Aggregately diversified recommender systems try to address this problem [1, 3, 6, 9, 10, 12]. However, these studies do not consider the order of recommendations and assume that all recommendations happen at once. In reality, the order of recommendations must be considered since users do not simultaneously receive recommendations. Thus, a recommender system cannot access the recommendation results for future users in advance.

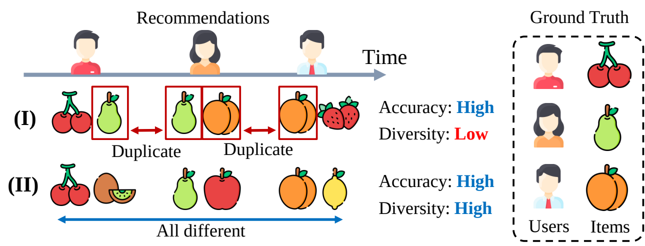

Figure 1: Examples of two sequential recommendation results (I) and (II).

We propose the problem of sequentially diversified recommendation to consider the order of recommendations. Figure 1 shows the example. Three users sequentially receive recommendations (I) and (II). Both results achieve the highest accuracy, but result (II) achieves higher diversity than result (I). This is because the result (I) recommends items that latter users would prefer, leading to duplicate recommendations. The sequentially diversified recommendation aims to maximize both accuracy and diversity as in the result (II). Thus, it is important to distinguish between the items to be recommended now and those to be recommended later.

In this paper, we propose SAPID (Sequentially Diversified Recommendation via Popularity Debiasing and Item Distribution), an effective method for the sequentially diversified recommendation. SAPID removes the popularity bias from the model by temporal popularity-based negative sampling scheme. Then, SAPID recommends items that would not have a chance to be recommended later to improve diversity. Extensive experiments in 4 real-world datasets show that SAPID achieves up to 4.25 times higher accuracy and up to 3.92% increased diversity compared to the best competitor. The code and datasets are available at https://github.com/SAPID24/SAPID.

2 PROBLEM FORMULATION

***Problem* 1** (Sequentially Diversified Recommendation)**:** Assume that a session is a sequence of user-item interactions for a single user with timestamps. For the set $\mathbb{U}$ of users, the sessions $\mathbb{S}_1, \cdots, \mathbb{S}_{|\mathbb{U}|}$ are given and sorted in their last interactions' time order. Each session is provided sequentially to perform a recommendation, one at a time. The problem is to recommend a list $\mathbb{R}_u$ of k items for each session $\mathbb{S}_u$ that are most likely to appear next in the session while maximizing both the accuracy and the aggregate-level diversity. □

Note that the recommender system needs to generate $\mathbb{R}_u$ as soon as $\mathbb{S}_u$ is provided. Thus, it is crucial to consider the future demand for items and decides which items to recommend immediately and which ones to reserve for each user to recommend to.

SIGKDD'24, August 25–29, 2024, Barcelona, Spain

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

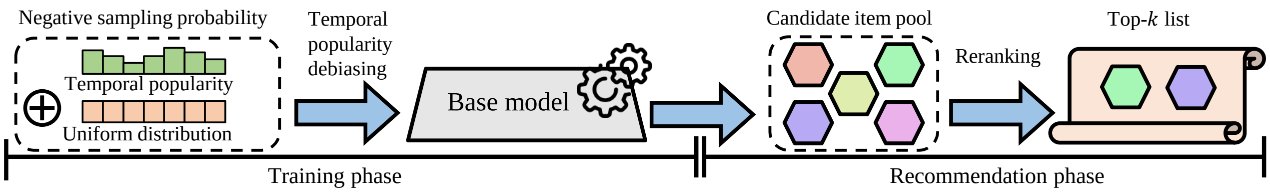

Figure 2: The overall architecture of SAPID.

3 PROPOSED METHOD

In this section, we propose SAPID (Sequentially Diversified Recommendation via Popularity Debiasing and Item Distribution), an effective approach for the sequentially diversified recommendation. Figure 2 shows the overall process of SAPID. In the training phase, SAPID carefully trains a sequential recommendation model not to be biased. Then, SAPID recommends items that contribute to the diversity the most in the recommendation phase.

Training Phase. How can we eliminate the popularity bias introduced by the skewed distribution during the model training process? The BPR loss in recommendation model training aims to maximize the difference between the recommendation scores of positive samples and randomly chosen negative samples [13]. However, this approach leads to the popularity bias in the model, as lesser-known items are underestimated, which results into the poor diversity of recommendation.

Our approach to address this issue is reducing the frequency of lesser-known items being chosen as negative samples to mitigate the popularity bias. SAPID calculates the temporal popularity $\mathbf{P}^{\mathbb{S}}_{pop}$ of negative items for a session $\mathbb{S}$ by counting all interactions of users and items not included in $\mathbb{S}$ that happened between the first interaction of the session and the last interaction of the session. Let $\mathbf{P}^{\mathbb{S}}_{uni}$ be the uniform distribution for all items not included in $\mathbb{S}$. Then, the probability distribution $\mathbf{P}^{\mathbb{S}}$ for each item to be selected as a negative sample for the session $\mathbb{S}$ is defined as Equation (1), where α is a hyperparameter which ranges from 0 to 1.

$$\mathbf{P}^{\mathbb{S}} = \alpha \mathbf{P}^{\mathbb{S}}_{pop} + (1-\alpha)\mathbf{P}^{\mathbb{S}}_{uni} \tag{1}$$

Recommendation Phase. How can we maintain the accuracy while increasing the diversity of recommendation results? SAPID utilizes the reranking method [1, 3, 6, 12] to maintain accuracy. SAPID first constructs the candidate item pool of size c to recommend by gathering highly ranked items, where c is a hyperparameter. Then, SAPID judges how much an item contributes to diversity by counting its expected number of being recommended. Finally, SAPID recommends items with the least expected number of being recommended from the candidate pool.

To calculate the expected number $E_u(i)$ of recommendations for item i calculated when recommending to the u'th user, SAPID sums 1) the number of times the item has been recommended for the previous sessions and 2) the estimated number of times it will be recommended for the remaining sessions. We assume that the recommendation rate of each item for future users is proportional to its popularity. Thus, $E_u(i)$ is defined as follows:

$$E_u(i) = C(i) + (|\mathbb{U}| - u)k \times p_i, \tag{2}$$

where $C(i)$ is the current count of recommendations for item i, and p_i is the frequency rate of item i.

Table 1: Summary of datasets.

Dataset	Users	Items	Interactions	Avg. Length
Gowalla	34,688	63,279	2,438,708	70.30
AMZ-elec	33,602	16,448	788,143	23.46
AMZ-home	12,000	8,445	291,560	24.30
Ml-1m	6,040	3,706	1,000,209	165.60

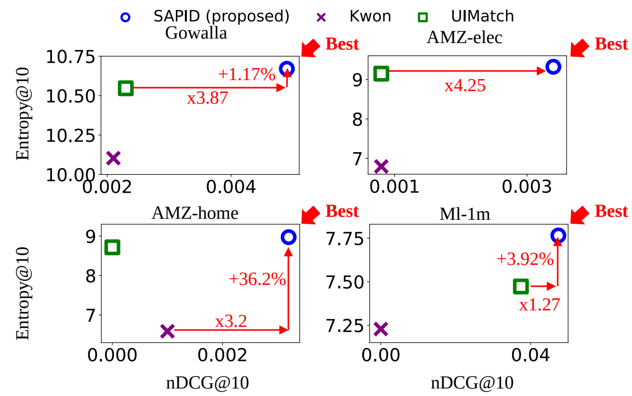

Figure 3: Performance of SAPID and competitors.

4 EXPERIMENT

We perform experiments to prove that SAPID achieve higher diversity while sacrificing lesser accuracy compared to competitors.

We use four real-world datasets of sequential recommendation as summarized in Table 1. We use SASRec [8], a widely used sequential recommendation models as our base models for experiments. We compare SAPID with two aggregately diversified recommendation methods: Kwon et al [1] and UImatch [3]. We employ a leave-one-out protocol and remove the last interaction of each user to construct the training data. We report 5-fold harmonic nDCG to measure the accuracy and entropy to measure the diversity.

Figure 3 shows the results. SAPID shows the best accuracy and diversity among competitors in all cases, which proves the effectiveness of SAPID for sequentially diversified recommendation.

5 CONCLUSION

In this paper, we propose the problem of sequentially diversified recommendation, an important problem to maximize the potential profit of e-commerce platforms. We also propose SAPID, an effective method for the problem. SAPID removes the popularity bias of the base model by adjusting the negative sampling probability. Then, SAPID decides which items to recommend immediately or later based on their potential chances to be exposed later. SAPID shows the state-of-the-art performance in real-world datasets by achieving up to 4.25 times higher accuracy and up to 3.92% increased diversity compared to the best competitor.

REFERENCES

[1] Gediminas Adomavicius and YoungOk Kwon. 2012. Improving Aggregate Recommendation Diversity Using Ranking-Based Techniques. *IEEE Transactions on Knowledge and Data Engineering* 24 (2012).

[2] Erik Brynjolfsson, Yu Hu, and Duncan Simester. 2011. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management science* 57, 8 (2011), 1373–1386.

[3] Qiang Dong, Shuang-Shuang Xie, and Wen-Jun Li. 2021. User-item matching for recommendation fairness. *IEEE Access* 9 (2021), 130389–130398.

[4] Daniel G Goldstein and Dominique C Goldstein. 2006. Profiting from the long tail. *Harvard Business Review* 84, 6 (2006), 24–28.

[5] Hyunsik Jeon, Jun-Gi Jang, Taehun Kim, and U Kang. 2023. Accurate bundle matching and generation via multitask learning with partially shared parameters. *Plos one* 18, 3 (2023), e0280630.

[6] Hyunsik Jeon, Jongjin Kim, Jaeri Lee, Jong-eun Lee, and U Kang. 2023. Aggregately diversified bundle recommendation via popularity debiasing and configuration-aware reranking. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 348–360.

[7] Hyunsik Jeon, Jongjin Kim, Hoyoung Yoon, Jaeri Lee, and U Kang. 2022. Accurate action recommendation for smart home via two-level encoders and commonsense knowledge. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 832–841.

[8] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.

[9] Mahmut Özge Karakaya and Tevfik Aytekin. 2018. Effective methods for increasing aggregate diversity in recommender systems. *knowledge and Information Systems* 56 (2018), 355–372.

[10] Jongjin Kim, Hyunsik Jeon, Jaeri Lee, and U Kang. 2023. Diversely regularized matrix factorization for accurate and aggregately diversified recommendation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 361–373.

[11] Bonhun Koo, Hyunsik Jeon, and U Kang. 2020. Accurate news recommendation coalescing personal and global temporal preferences. In *Advances in Knowledge Discovery and Data Mining: 24th Pacific-Asia Conference, PAKDD 2020, Singapore, May 11–14, 2020, Proceedings, Part I 24*. Springer, 78–90.

[12] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. FairMatch: A Graph-based Approach for Improving Aggregate Diversity in Recommender Systems. In *UMAP 2020 - Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. Association for Computing Machinery, Inc.

[13] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).