Accurate and Diverse Recommendation based on Users' Tendencies toward Temporal Item Popularity

Koki Nagatani
Fuji Xerox Co., Ltd.
6-1 Minatomirai, Nishi-ku, Yokohama, Japan
nagatani.koki@fujixerox.co.jp

Masahiro Sato
Fuji Xerox Co., Ltd.
6-1 Minatomirai, Nishi-ku, Yokohama, Japan
sato.masahiro@fujixerox.co.jp

ABSTRACT

Popularity bias is a phenomenon associated with collaborative filtering algorithms, in which popular items tend to be recommended over unpopular items. As the appropriate level of item popularity differs depending on individual users, a user-level modification approach can produce diverse recommendations while improving the recommendation accuracy. However, there are two issues with conventional user-level approaches. First, these approaches do not isolate users' preferences from their tendencies toward item popularity clearly. Second, they do not consider temporal item popularity, although item popularity changes dynamically over time in reality. In this paper, we propose a novel approach to counteract the popularity bias, namely, matrix factorization based collaborative filtering incorporating individual users' tendencies toward item popularity. Our model clearly isolates users' preferences from their tendencies toward popularity. In addition, we consider the temporal item popularity and incorporate it into our model. Experimental results using a real-world dataset show that our model improve both accuracy and diversity compared with a baseline algorithm in both static and time-varying models. Moreover, our model outperforms conventional approaches in terms of accuracy with the same diversity level. Furthermore, we show that our proposed model recommends items by capturing users' tendencies toward item popularity: it recommends popular items for the user who likes popular items, while recommending unpopular items for those who don't like popular items.

CCS CONCEPTS

• **Information systems** → **Personalization**; **Recommender systems**;

KEYWORDS

popularity bias, temporal information, personalized recommendation

1 INTRODUCTION

Recommender systems help users to access the specific information that they seek from a huge amount of data. Accurate recommendations lead to an increase in customers' purchases or consumption; hence, there is a need for more efficient recommender systems that produce personalized content for individual users. To produce personalized recommendations, collaborative filtering (CF) is a widely used approach. The CF approach produces items for a target user using data compiled from observations of users with similar preferences as the target user [9]. The CF approach is categorized into two types: neighborhood-based CF [12, 13] and model-based CF [5, 11]. The standard approach of model-based CF is a matrix factorization (MF)-based approach, which characterizes both items and users by vectors of latent factors inferred from user feedback [5, 11]. In most cases, model-based CF is superior to neighborhood-based CF in terms of accuracy.

In the CF approach, it has been noted that popular items tend to be recommended more often [15, 21]. This is known as popularity bias and various solutions have been proposed to tackle this problem [3, 4, 7, 10, 21]. These solutions are classified into two types according to the level of modification: global-level and user-level. Global-level solutions modify their recommendations for all users uniformly by avoiding recommending popular items [3, 7, 21]. In reality, however, the appropriate level of modification differs depending on the user: some users are likely to select popular items, while others tend to seek new or niche items. Therefore, user-level modification approaches, in which the degree of modification varies according to individual users' popularity tendencies, have been proposed [4, 10].

However, there are two issues in conventional user-level approaches. First, these approaches do not isolate users' preferences from their popularity tendencies clearly. Second, although item popularity changes dynamically over time in reality, these approaches do not consider temporal item popularity. In general, incorporating temporal item popularity into models improves the recommendation accuracy. Moreover, to counteract popularity bias, especially in user-level solutions, incorporating temporal item popularity is important because the reasons of users' behaviors are considered different depending on their purchase time even if they purchase same items. To the best of our knowledge, however, there is no approach considering temporal item popularity in the field of counteraction against popularity bias.

In this paper, we propose a novel approach to tackle the popularity bias, namely, MF-based CF incorporating item popularity orientation of individual users. Our model isolates users' preference from their tendencies toward item popularity clearly. We also consider temporal item popularity and incorporate it into our model. To verify the efficacy of the proposed model, we conducted experiments using a real-world dataset. The experimental results show that our model improves both accuracy and diversity compared with a baseline algorithm in both static and time-changing models. Moreover, our model outperforms conventional approaches in

TempRRS '17, August 2017, Como, Italy

terms of accuracy with the same diversity level. We also demonstrate that our proposed model recommends items by capturing users' tendencies toward item popularity: it recommends popular items to users who like popular items, and unpopular items to those who do not like popular items.

We summarize the main contribution of this paper as follows:

- Our model isolates users' preferences from their tendencies toward item popularity clearly.
- We consider temporal item popularity in the field of counteraction against popularity bias.
- We conduct experiments using a real-world dataset to verify the efficacy of the proposed model.

2 RELATED WORK

2.1 Popularity Bias

Popularity bias is a phenomenon of existing recommendation algorithms in which popular items tend to be recommended over unpopular items. To tackle this problem, several approaches have been proposed [3, 4, 7, 10, 21]. These approaches are classified into two types according to the level of modification: global-level and user-level.

Global-level approaches modify their recommendations for all users uniformly [3, 7, 21]. Most methods avoid recommending popular items by weighting according to item popularity. In global-level approaches, the evaluation metrics such as diversity and novelty improve at the cost of a decline in accuracy. Generally, the appropriate level of modification differs depending on the user: some users are likely to select popular items, some tend to seek new or niche items, and some select an item irrespective of its popularity. However, global-level approaches do not consider such individual differences.

User-level approaches consider these differences and then modify their recommendations depending on the individual user's tendencies toward item popularity. Therefore, user-level approaches possibly improve both diversity and accuracy simultaneously. The conventional user-level approaches proposed in [4, 10] attempt to re-rank recommendation lists by post sampling based on users' past behavior in terms of popularity. However, users' preferences and their tendencies toward item popularity might be mixed in these approaches for two reasons. First, before reranking, the recommendation lists are created by existing CF models. During the creation process, these models mix users' preference and item popularity. Second, popularity tendency distributions are created based on users' past actions. As users' past actions are mainly derived from the users' preferences and items' popularity, these aspects are also included when creating the distribution. Therefore, these approaches do not isolate user preferences from their popularity tendencies clearly. Our solution overcomes the above issue by modeling users' popularity tendencies directly, as described in Section 3.

2.2 CF with Temporal Aspects

Incorporating temporal aspects into CF has been investigated, particularly for developing accurate recommendation algorithms. For example, [6] proposed a matrix factorization model that considered temporal dynamics and achieved state-of-the-art performance at the time on Netflix data. Since then, several models that consider temporal dynamics using MF [1, 20] or deep learning methods [16, 19] have been proposed. In user-level approaches to popularity bias, temporal item popularity needs to be considered to capture individual users' tendencies toward item popularity. This is because the reasons for purchasing items in case of users having multiple interactions with the same items may be different depending on the interaction time: some users purchase items because the items are popular, and some users purchase items because the items match the users' preferences. To our knowledge, however, there is no approach that considers temporal aspects in the field of popularity bias.

3 OUR MODEL

In this section, we present our MF-based model that incorporates individual users' tendencies toward item popularity. We focus on situations where personalized top-N recommendations are produced based on users' implicit feedback (e.g. views, clicks, purchases, etc.).

3.1 Modeling Individual Users' Tendencies toward Temporal Item Popularity

In MF, both items and users are characterized by vectors of latent factors derived from explicit feedback (e.g. ratings) as well as implicit feedback. The basic model of MF with item bias is formulated as follows:

$$\hat{x}_{ui} = b_i^0 + f_u^T f_i, \tag{1}$$

where $\hat{x}_{ui}$ is the prediction score of preference of user u toward item i, b_i^0 is an item-specific bias which represents item popularity, and f_u and f_i are k-dimensional vectors of latent factors of user u and item i, respectively. The inner product $f_u^T f_i$ achieves a high value when both user and item vectors are similar. Furthermore, item bias b_i increases when an item is popular. The prediction score is determined by their aggregation.

If item bias b_i values are extremely high, the item is recommended regardless of whether users like it or not. Hence, recommendation systems tend to recommend these items, which leads to popularity biased recommendation. A simple solution for this problem is to penalize items according to the item popularity. However, preference toward popular or unpopular items varies for each user. Considering this, the solution is not suitable for users who like popular items. Therefore, the penalization of popularity needs to be changed depending on the users' popularity tendencies.

Moreover, the users' popularity tendencies should be considered along with the items' temporal aspects for two reasons. Firstly, item popularity changes dynamically over time in the real world for various reasons [17]. Secondly, the reasons for purchasing items in case of users having multiple interactions with the same items may be different depending on the interaction time.

Therefore, we develop a model to incorporate both users' popularity tendencies and items' temporal popularity, which is formulated as follows:

$$\hat{x}_{ui} = (b_i^0 + b_i(t))(1 + g_u) + f_u^T f_i, \tag{2}$$

where g_u is the user-specific parameter of popularity tendency and $b_i(t)$ is the time-varying item bias at the period of time t. The parameters, b_i^0, $b_i(t)$, g_u, f_u, and f_i, are learned by optimization.

The g_u value works as the balancing parameter between the item popularity and preference toward the item. When the g_u value of a user u is greater than zero, the user prefers popular items to unpopular items. High g_u values indicated that the user may simply prefer popular items without regard to his/her item preference. Conversely, when it is less than minus one, the user prefers unpopular items to popular items.

As mentioned in Section 2.1, users' preferences and their tendencies toward item popularity are mixed in conventional user-level approaches. In contrast, our model resolves the confusion by modeling as in Eq. 2: the first term represents item popularity and users' popularity tendencies, and the second term represents item feature and users' preference. Therefore, our model captures these features separately.

3.2 Model Learning

Our model formulated in Eq. 2 can be learned by applying existing optimization methods, such as point-wise and pair-wise optimization. For example, for point-wise optimization, root mean square error (RMSE), which is used in Biased-MF [5], and alternating least squares, which is used in weighted regularize matrix factorization [2] can be applied to our model. For pair-wise optimization, area under the curve (AUC) in Bayesian personalized ranking [11], mean reciprocal rank used in collaborative less-is-more filtering (CLiMF) [14], and weighted approximately ranked pairwise loss proposed in [18] can be applied to our models.

4 EXPERIMENTS

In this section, we conduct experiments using a real-world dataset to verify the efficacy of the proposed model.

4.1 Dataset

We used the Amazon.com Movies and TVs dataset [8] in our experiment. We utilized a subset from 2013, defined the period of time t as monthly, and binarized the data treating reviewed items as relevant and non-reviewed items as irrelevant. Due to the sparsity of the dataset, we preprocessed it by retaining the top $10,000$ items and discarding data of users having less than 10 interactions. After the preprocessing, the total number of users was $4,997$ and the dataset contained $90,341$ interactions for $9,221$ items.

4.2 Evaluation Metrics

In our experiments, we performed five-fold cross validation and aggregated the results. First, we randomly selected 80% of observed feedback as a training set to train models, and the remaining 20% as the testing set for the trained models. To measure the performance, we used three evaluation metrics: the top-N prediction precision (Precision@N), the top-N prediction recall (Recall@N), and the top-N item coverage (Coverage@N). We set $I_u^{\text{pred}}(t)$ as the predicted items of user u over a certain period of time t, and $I_u^{\text{true}}(t)$ as the true list in the testing set. Prediction is performed for each period of time t, and each user's scores are aggregated over each period. The top-N prediction precision is defined as:

$$\text{Precision@}N = \frac{1}{|U|} \sum_{u \in U} \frac{1}{|T_u|} \sum_{t \in T_u} \frac{|I_u^{\text{pred}}(t) \cap I_u^{\text{true}}(t)|}{N},$$

where $|I_u^{\text{pred}}(t)| = N$, U is the set of users in the testing set and T_u is the set of the period of time when interactions of user u are observed. Similarly, the top-N prediction recall is defined as:

$$\text{Recall@}N = \frac{1}{|U|} \sum_{u \in U} \frac{1}{|T_u|} \sum_{t \in T_u} \frac{|I_u^{\text{pred}}(t) \cap I_u^{\text{true}}(t)|}{|I_u^{\text{true}}(t)|}.$$

The top-N item coverage applies to all the output that a recommender system produces for a set of users. This metric is also called the top-N aggregate diversity. In our experiment, this metric is defined as:

$$\text{Coverage@}N = \frac{\sum_{t \in T} |\bigcup_{u \in U_t} R_u|}{|T|},$$

where U_t is a set of users whose interactions are observed at a period of time t and R_u is the recommendation lists for user u, and the length of the lists is N.

4.3 Comparison of Methods

To examine the performance of our proposed methods, we compared them with conventional approaches. For the optimization of our methods and base models of conventional approaches, we selected the Bayesian personalized ranking (BPR) procedure [11], which is one of the state-of-the-art methods for personalized item recommendation. The model of BPR matrix factorization (BPRMF) is formulated in Eq. 1. For the baseline of the conventional methods that consider temporal aspects, we extend BPRMF incorporating temporal item popularity, which is called BPRMF(t).

Personal Popularity Tendency Matching (PPTM) [10] is a greedy re-ranking method that considers an individual's personal popularity tendency (PPT). It balances novelty and user preference by matching the PPT of a recommendation to that of the users measured by earth movers distance (EMD), which is a distance metrics between two distributions.

Personalized Ranking Adaptation (PRA) [4] is a versatile greedy re-ranking method that considers an individual user tendency suitable for multiple optimization goals. In our experiments, the optimization target is set to EMD.

BPRMF(t)-pop is the method proposed by this paper in Eq. 2. BPRMF-pop is the model that removes temporal item popularity from Eq. 2.

To model PPT, the discrete distribution of the binned popularity values of the items is required. In our experiments, we defined the item popularity of the recommendations as the number of item occurrences in the top-N recommendation lists for the active users. We used a log-scaled popularity histogram for discrete distribution. The parameters of all models were tuned so as to maximize the accuracy metrics. In the case of conventional approaches, it is known that a higher coverage setting reduces accuracy. Hence, we selected the parameter value for which the coverage score became close to that of our model.

Table 1: Precision@10, Recall@10, and Coverage@10 scores on Amazon.com Movies and TVs datasets (#factors = 300).

Methods	Precision@10	Recall@10	Coverage@10
BPRMF	0.01349	0.07492	0.6464
+PPTM ($c = 0.1$)	0.01348	0.07487	0.6509
+PRA (X_u=5)	0.01274	0.07012	0.7029
BPRMF-pop	0.01359	0.07550	0.6528
BPRMF(t)	0.01521	0.09314	0.4939
+PPTM ($c = 1$)	0.01504	0.09047	0.5634
+PRA (X_u=5)	0.01308	0.07764	0.5605
BPRMF(t)-pop	0.01603	0.09569	0.5749

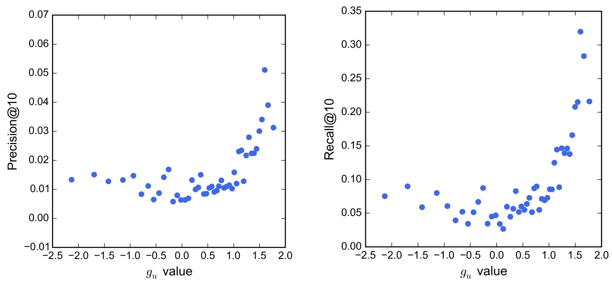

(a) g_u vs. Precision@10 (b) g_u vs. Recall@10

Figure 1: Plots of g_u values versus Precision@10 and Recall@10. Each point is the average of evaluation metrics with regards to the average of g_u values of 100 users in descending order of the g_u value.

4.4 Experimental Results

Table 1 shows the results of the comparison between our method and conventional approaches. The number of latent factors was set to 300 and the number of items in a recommendation list to 10. In general, time-aware models improve accuracy and reduce coverage compared with static models. Our model improved both accuracy and diversity compared with the baseline in both static and time-changing models. Particularly in case of time-varying models, our model achieved significant improvement. This indicates that considering temporal item popularity is essential to capture users' tendencies. Our model outperformed conventional approaches in terms of accuracy with the same diversity level. Therefore, our model effectively captures users' preference and their tendencies toward item popularity.

We suppose that the interactions of users with mainstream tastes are easy to predict. As our model isolates users' preference from their tendencies toward item popularity, we can verify the idea by analyzing the distribution of accuracy depending on the magnitude of g_u values. Figure 1 shows the plots of g_u values versus two evaluation metrics: Precision@10 and Recall@10; each point is the average of evaluation metrics with regards to the average of g_u values of 100 users in descending order of the g_u value. As can be seen from Fig. 1, both Precision@10 and Recall@10 of the users who have large g_u values are high. Therefore, this result supports our assumption.

We also investigated the relation between users' purchase behavior, which corresponds to their tendencies toward temporal item popularity, and our model's recommendations. Table 2 shows

Table 2: The $b_i(t)$ and preference score of Top-5 Recommendation of BPRMF(t)-pop for a user at time t. Actual user behavior (relevant recommendation) is bolded. Ranking is calculated based on the training set. "-" means the item is not in the training set.

(a) Item popularity orientation score $g_u = 0.69$, a user likes popular items.

TopN items	$t = 3$		$t = 12$	
	$b_i(t)$ (#rank)	Pref. score	$b_i(t)$ (#rank)	Pref. score
1	4.17 (1)	0.06	4.63 (4)	0.041
2	3.95 (3)	0.02	4.33 (1)	-0.017
3	3.73 (4)	0.02	4.14 (3)	-0.017
4	**3.62 (14)**	**0.10**	4.10 (6)	-0.028
5	3.67 (8)	-0.02	**4.04 (7)**	**0.012**

(b) Item popularity orientation score $g_u = -1.19$,
a user selects items which match user's preferences.

TopN items	$t = 2$		$t = 3$	
	$b_i(t)$ (#rank)	Pref. score	$b_i(t)$ (#rank)	Pref. score
1	0.28 (1458)	2.29	-1.16 (1123)	2.60
2	-0.44 (1458)	1.98	1.70 (169)	2.83
3	**-0.72 (-)**	**1.76**	-1.13 (-)	1.69
4	-2.13 (1458)	1.47	1.75 (492)	2.17
5	0.36 (1458)	1.90	**0.25 (492)**	**1.88**

the examples that our model recommends popular items for the user who likes popular items and vice versa. $b_i(t)$ score represents item popularity at the period of time t and actual users' purchase is shown in bold in Table 2. As can be seen from the user's purchase behavior shown in Table 2-(a), the user tends to purchase popular items. Our model learned such purchase behavior from the user's past purchases, and then evaluated the g_u value of the user as 0.69, which means that the user likes popular items. Our model produced popular items for the user, which were ranked in the top 20. On the other hand, the user in Table 2-(b) selected items that match the user's preference without regard to the items' popularity. Our model captured the tendency from the user's past purchases and evaluated the user's g_u value as -1.19. Our model recommended items that match the user's preference regardless of their popularity for the user. The preference scores are all high, while the items' rankings are various. Therefore, these results indicate that our model captured the users' popularity tendencies and recommended personalized items appropriately.

5 CONCLUSIONS

In this paper, we proposed a novel approach for counteracting popularity bias, using MF-based CF incorporating individual users' tendencies toward temporal item popularity. Our model isolated users' preference from popularity tendency clearly, and considered temporal item popularity. The experimental results based on a real-world dataset showed the efficacy of our model.

In future work, we plan to further verify the effectiveness of our proposed model by using various datasets in different domains or by learning other optimization methods for top-N recommendation. Moreover, as well as item popularity, users' tendencies toward item popularity may change over time. We plan to investigate this temporal phenomenon.

REFERENCES

[1] Ruining He, Chen Fang, Zhaowen Wang, and Julian McAuley. 2016. Vista: A Visually, Socially, and Temporally-aware Model for Artistic Recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 309–316. https://doi.org/10.1145/2959100.2959152

[2] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining (ICDM '08)*. IEEE Computer Society, Washington, DC, USA, 263–272. https://doi.org/10.1109/ICDM.2008.22

[3] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What Recommenders Recommend: An Analysis of Recommendation Biases and Possible Countermeasures. *User Modeling and User-Adapted Interaction* 25, 5 (Dec. 2015), 427–491. https://doi.org/10.1007/s11257-015-9165-3

[4] Michael Jugovac, Dietmar Jannach, and Lukas Lerche. 2017. Efficient Optimization of Multiple Recommendation Quality Factors According to Individual User Tendencies. *Expert Syst. Appl.* 81, C (Sept. 2017), 321–331. https://doi.org/10.1016/j.eswa.2017.03.055

[5] Yehuda Koren. 2008. Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '08)*. ACM, New York, NY, USA, 426–434. https://doi.org/10.1145/1401890.1401944

[6] Yehuda Koren. 2009. Collaborative Filtering with Temporal Dynamics. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '09)*. ACM, New York, NY, USA, 447–456. https://doi.org/10.1145/1557019.1557072

[7] Qiuxia Lu, Tianqi Chen, Weinan Zhang, Diyi Yang, and Yong Yu. 2012. Serendipitous Personalized Ranking for Top-N Recommendation. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology - Volume 01 (WI-IAT '12)*. IEEE Computer Society, Washington, DC, USA, 258–265. http://dl.acm.org/citation.cfm?id=2457524.2457692

[8] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-Based Recommendations on Styles and Substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*. ACM, New York, NY, USA, 43–52. https://doi.org/10.1145/2766462.2767755

[9] Xia Ning, Christian Desrosiers, and George Karypis. 2015. A Comprehensive Survey of Neighborhood-Based Recommendation Methods, In *Recommender Systems Handbook*. Springer, New York City, New York.

[10] Jinoh Oh, Sun Park, Hwanjo Yu, Min Song, and Seung-Taek Park. 2011. Novel Recommendation Based on Personal Popularity Tendency. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining (ICDM '11)*. IEEE Computer Society, Washington, DC, USA, 507–516. https://doi.org/10.1109/ICDM.2011.110

[11] 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 (UAI '09)*. AUAI Press, Arlington, Virginia, United States, 452–461. http://dl.acm.org/citation.cfm?id=1795114.1795167

[12] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2000. Analysis of Recommendation Algorithms for e-Commerce. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce (EC '00)*. ACM, New York, NY, USA, 158–167. https://doi.org/10.1145/352871.352887

[13] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In *Proceedings of the 10th International Conference on World Wide Web (WWW '01)*. ACM, New York, NY, USA, 285–295. https://doi.org/10.1145/371920.372071

[14] Yue Shi, Alexandros Karatzoglou, Linas Baltrunas, Martha Larson, Nuria Oliver, and Alan Hanjalic. 2012. CLiMF: Learning to Maximize Reciprocal Rank with Collaborative Less-is-more Filtering. In *Proceedings of the Sixth ACM Conference on Recommender Systems (RecSys '12)*. ACM, New York, NY, USA, 139–146. https://doi.org/10.1145/2365952.2365981

[15] Harald Steck. 2010. Training and Testing of Recommender Systems on Data Missing Not at Random. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '10)*. ACM, New York, NY, USA, 713–722. https://doi.org/10.1145/1835804.1835895

[16] Yong Kiam Tan, Xinxing Xu, and Yong Liu. 2016. Improved Recurrent Neural Networks for Session-based Recommendations. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems (DLRS 2016)*. ACM, New York, NY, USA, 17–22. https://doi.org/10.1145/2988450.2988452

[17] Yusuke Tanaka, Takeshi Kurashima, Yasuhiro Fujiwara, Tomoharu Iwata, and Hiroshi Sawada. 2016. Inferring Latent Triggers of Purchases with Consideration of Social Effects and Media Advertisements. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM '16)*. ACM, New York, NY, USA, 543–552. https://doi.org/10.1145/2835776.2835789

[18] Jason Weston, Hector Yee, and Ron J. Weiss. 2013. Learning to Rank Recommendations with the K-order Statistic Loss. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, New York, NY, USA, 245–248. https://doi.org/10.1145/2507157.2507210

[19] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J. Smola, and How Jing. 2017. Recurrent Recommender Networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*. ACM, New York, NY, USA, 495–503. https://doi.org/10.1145/3018661.3018689

[20] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. 2016. Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 847–855.

[21] Xiangyu Zhao, Zhendong Niu, and Wei Chen. 2013. Opinion-based Collaborative Filtering to Solve Popularity Bias in Recommneder Systems. In *24th International Conference on Database and Expert Systems Applications (DEXA)*. 426–433.