

Dynamic-K Recommendation with Personalized Decision Boundary

Yan Gao, Jiafeng Guo, Yanyan Lan, and Huaming Liao

CAS Key Lab of Network Data Science and Technology Institute of Computing
Technology, Chinese Academy of Sciences
{gaoyan}@software.ict.ac.cn,
{guojiafeng,lanyanyan,lhm}@ict.ac.cn

Abstract. In this paper, we investigate the recommendation task in the most common scenario with implicit feedback (e.g., clicks, purchases). State-of-the-art methods in this direction usually cast the problem as to learn a personalized ranking on a set of items (e.g., webpages, products). The top-N results are then provided to users as recommendations, where the N is usually a fixed number pre-defined by the system according to some heuristic criteria (e.g., page size, screen size). There is one major assumption underlying this fixed-number recommendation scheme, i.e., there are always sufficient relevant items to users' preferences. Unfortunately, this assumption may not always hold in real-world scenarios. In some applications, there might be very limited candidate items to recommend, and some users may have very high relevance requirement in recommendation. In this way, even the top-1 ranked item may not be relevant to a user's preference. Therefore, we argue that it is critical to provide a dynamic-K recommendation, where the K should be different with respect to the candidate item set and the target user. We formulate this dynamic-K recommendation task as a joint learning problem with both ranking and classification objectives. The ranking objective is the same as existing methods, i.e., to create a ranking list of items according to users' interests. The classification objective is unique in this work, which aims to learn a personalized decision boundary to differentiate the relevant items from irrelevant items. Based on these ideas, we extend two state-of-the-art ranking-based recommendation methods, i.e., BPRMF and HRM, to the corresponding dynamic-K versions, namely DK-BPRMF and DK-HRM. Our experimental results on two datasets show that the dynamic-K models are more effective than the original fixed-N recommendation methods.

Keywords: Implicit Feedback, Dynamic-K Recommendation

1 Introduction

Recommender systems have been widely used in many applications, such as Amazon, YouTube and so on. In this paper, we address the most common recommendation scenario with implicit feedbacks, e.g., clicks or purchases from

users. Most methods in this direction cast the problem as to learn a personalized ranking on a set of items[6], e.g., webpages or products. The top-N ranked items are then provided to users as recommendations[1,8,11], where N is usually a fixed number pre-defined by the system according to the recommendation space (e.g., page size, screen size) or some heuristic criteria.

There is an underlying assumption for such fixed-N recommendation scheme, i.e., there are always sufficient relevant items to users' preferences. Unfortunately, this assumption may not always hold in real-world scenarios. For example, if one aims to recommend newly uploaded papers on arxiv¹ for academic readers everyday, he/she may face the following two problems. Firstly, there might be very limited new papers updated on arxiv everyday. Secondly, some academic readers might have very high relevance requirement in recommendation since paper reading is time-consuming. In this way, even the top-1 ranked paper may not be appropriate to be recommended to such readers on some day. However, by using a fixed-N recommendation scheme, each reader will constantly receive N recommended papers, and it is very likely to annoy the readers in this situation.

To avoid the above problem, we argue that it is critical to provide a dynamic-K recommendation, where the K should be different with respect to the candidate item set and the target user. If there are sufficient relevant items or if the user likes receiving diverse recommendations (i.e., relatively low relevance requirement), the K could be large. In contrary, if there are limited relevant items or if the user likes receiving recommendations unless they are highly relevant, the K should be small or even no recommendation should be provided sometimes. Ideally, a good recommender system should be able to learn from users' implicit feedbacks to present such dynamic-K recommendations for different candidate item set and target users.

In this work, we formulate this dynamic-K recommendation task as a joint learning problem with both ranking and classification objectives. Specifically, the ranking objective is the same as many existing ranking-based methods, i.e., to create a ranking list of items according to users' interests. The classification objective, which is unique in this work, aims to learn a personalized rather than a global decision boundary to differentiate the relevant items from irrelevant items. We apply the above joint learning idea over two state-of-the-art ranking-based recommendation methods, i.e., BPRMF and HRM, and extend them to the corresponding dynamic-K versions, namely DK-BPRMF and DK-HRM, respectively.

We conduct empirical experiments over two publicly available datasets, i.e., a transaction dataset named Ta-Feng and a movie recommendation dataset named MovieLens. Our experimental results show that the dynamic-K models are more effective than the corresponding fixed-N recommendation methods as well as those existing hybrid recommendation methods.

¹ <https://arxiv.org/>

2 Related Work

Many methods have been developed in literature to build implicit feedback recommendation systems. With respect to the objective function, the recommendation methods can be further divided into three folds, namely ranking-based methods, classification-based methods and hybrid methods.

Ranking-based recommendation methods, which aim to correctly rank items rather than to correctly predict their ratings have demonstrated good performance for top-N recommendation system. Ranking-based recommendation methods can be categorized into two folds, i.e., pair-wise approach and list-wise approach. Pair-wise approach aims to optimize the pair-wise loss. For example, Rendle et al.[10] and Aioli et al.[2] optimized AUC score; Yun et al.[18] explored the connection between the Discounted Cumulative Gain(DCG) and the binary classification to change the ranking problem into binary classification problems; Park et al.[9] proposed a large-scale collaborative ranking method to minimize the ranking risk in the reconstructed recommendation matrix. List-wise approach optimizes the preference of each user to a list of items. An important branch in this category is designed to directly optimize evaluation metrics, such as Mean Average Precision(MAP), Mean Reciprocal Rank(MRR) and Normalized Discounted Cumulative Gain(NDCG), which are usually list-wise ranking metrics. Typical methods include TFMAP[14], CLiMF[15] and CofiRank[17]. Different types of ranking approaches have different strengths in producing the ranking list, but all these ranking-based methods are proposed under the fixed-N recommendation scenario.

Classification-based recommendation methods attempt to predict whether a user would like to interact with an item. Mnih et al.[7] provide a probabilistic framework for the implicit case where they model the probability of a user choosing an item according to a normalized exponential function. They avoid linear time computation and approximate the normalization to the distribution by traversing a tree structure. Goplan et al.[4] introduced a factorization model that factorizes users and items by Poisson distribution. More recently, Johnson [5] proposed a new probabilistic framework for the implicit case in which they model the probability of a user choosing an item by a logistic function. Nevertheless, these methods do not learn a personalized classifier for each user, thus can't be used for dynamic-K recommendation scenario.

Hybrid recommendation methods try to optimize several objectives simultaneously which can integrate the complementary strengths of different types of losses. Zhao et al.[3] propose a personalized top-N recommendation approach that minimizes a combined heterogeneous loss based on linear self-recovery models. The heterogeneous loss integrates the strengths of both pair-wise ranking loss and point-wise recovery loss to provide more informative recommendation predictions. Compared with their heterogeneous loss based on linear self-recovery models, we formulate our dynamic-K recommendation task as a joint learning problem with both ranking and classification objectives. Specifically, our classification objective aims to learn a personalized decision boundary for each user to differentiate the relevant items from irrelevant items.

3 Our Approach

In this section, we first introduce the problem formalization of recommendation with implicit feedback. We then describe the joint learning approach in detail. After that, we show how it can be applied to extend two state-of-the-art ranking-based recommendation methods, i.e., BPRMF and HRM, to the corresponding dynamic-K versions, namely DK-BPRMF and DK-HRM.

3.1 Problem Definition

Let U be a set of users and I be a set of items. In our scenario, implicit feedback $X \subseteq U \times V$ is available. Each instance $(u, i) \in X$ is a pair which means an interaction between user u and item i . Formally, we define the set of items with which user u has interactions is $B_u^+ := \{i \in I | (u, i) \in X\}$. With the learned model, we generate a list of all candidate items for each user. The top K ranked items are then provided to users as recommendations, where K is different with respect to the candidate item set and the target user. That is called dynamic- K recommendation system.

3.2 Joint Learning Approach

In this work, we formulate the dynamic- K recommendation task as a joint learning problem with both ranking and classification objectives. Specifically, the ranking objective could use any existing ranking-based methods, i.e., to create a ranking list of items according to user’s interests. Since we are dealing with implicit feedback data, we follow the rationale from Rendle et al.[10] assuming that user u prefers item i over item j if $i \in B_u^+ \wedge j \notin B_u^+$. Formally, we define the set P as the set of tuples $\{(u, i, j)\}$ selected from dataset X as follows: $P = \{(u, i, j) | i \in B_u^+ \wedge j \notin B_u^+\}$. Therefore, the ranking objective can be formalized as:

$$L(P; \theta)_{rank} = \sum_{(u, i, j) \in P} \ell(s_{u, i} \succ s_{u, j}; \theta) \quad (1)$$

where θ represents the parameter vector of an arbitrary model(e.g. matrix factorization). $s_{u, i}$, $s_{u, j}$ means the predicted score of item i to user u and item j to user u respectively. $\ell(\cdot)$ can be any arbitrary ranking loss function.

The classification objective, which is unique in this work, aims to learn a personalized decision boundary to differentiate the relevant items from irrelevant items. Intuitively, if the predicted score of the item is above user’s decision boundary, the item is relevant. If the predicted score is under user’s decision boundary, the item is irrelevant. We use a variable called $margin_{ui}$ to measure the confidence of the predicted class of item i to user u (i.e., if item i is relevant or irrelevant to user u):

$$margin_{ui} = y_{ui}(s_{u, i} - t_u) \quad (2)$$

where $y_{ui} \in \{-1, 1\}$ is the target class, t_u is the personalized decision boundary for each user. $s_{u,i}$ is the predicted score of item i to user u with an arbitrary model (e.g. matrix factorization). $Margin_{ui} < 0$ indicates item i is misclassified to user u , while $margin_{ui} > 0$ indicates item i is correctly classified to user u . $Margin_{ui}$ represents the "margin of safety" by which the prediction for item i to user u is correct. Furthermore, we assume the personalized decision boundary is regularized by a global constraint to avoid over-fitting. For implicit feedback data, we define the dataset D as a set of tuples $\{(u, i, y_{ui})\}$ selected from $U \times I$, and y_{ui} is defined as follows:

$$y_{ui} = \begin{cases} 1 & (u, i) \in X \\ -1 & (u, i) \notin X \end{cases} \quad (3)$$

Therefore, the learning objective in classification is to find the best θ and t_u to minimize classification loss over the training data D as following:

$$L(D; \theta, t_u)_{cf} = \sum_{(u, i, y_{ui}) \in D} \ell(margin_{ui}; \theta) + \lambda_t R(t_u - t) \quad (4)$$

where $\ell(\cdot)$ is defined as a loss function of the margin for each data, t is the global constrain we set for t_u and $R(\cdot)$ is a regularizer (typically L2 or L1). $\lambda_t > 0$ is a coefficient controlling the regularization strength. As noted above, any classification loss function could be applied in this framework. In particular, we explored the use of logistic loss and hinge loss in preliminary experiments, but found that logistic loss performed better on the data sets in this paper. Logistic loss is as following:

$$\ell_{log} = \ln(1 + e^{-margin_{ui}}) \quad (5)$$

Finally, by combining (1) and (4), we obtain our joint learning approach as follows:

$$L_{hybrid} = \alpha L_{cf} + (1 - \alpha) L_{rk} \quad (6)$$

where the parameter $\alpha \in [1, 0]$ denotes the trade-off between optimizing ranking loss and classification loss. Note that by setting $\alpha = 1$, the model reduces to a classification-based model; By setting $\alpha = 0$, we obtain a ranking-based model.

A direct algorithm for optimizing the joint objective function would enumerate the full set P of candidate pairs. Because $|P|$ is quadratic in $|D|$, this would be intractable for large-scale data sets. Instead, following similar idea with Sculley[12], we take the sampling approach from P rather than constructing P explicitly. Algorithm 1 gives a method for efficiently solving the joint learning optimization problem using stochastic gradient descent (SGD)[13].

In prediction stage, with the learned model, the dynamic-K recommendation with our joint learning approach is as following. Given a user u and an itemset I , for each candidate item $i \in I$, we calculate the predicted score s_{ui} . We then rank the items according to their predicted score, and select the top K results above the threshold t_u as the final recommendation results.

Algorithm 1 Joint Learning Approach. Given: tradeoff parameter α , regularization parameter λ_t , iterations N

```

1:  $\theta, t_u \leftarrow \emptyset, t$ 
2: for  $n = 1$  to  $N$  do
3:   pick  $z$  uniformly at random from  $[0, 1]$ 
4:   if  $z < \alpha$  then
5:      $(u, i, y_{ui}) \leftarrow \text{RandomExample}(D)$ 
6:      $L_{cf} = \ell(\text{margin}_{ui}) + \lambda_t R(t_u - t)$ 
7:      $\theta^n \leftarrow \text{StochasticGradientStep}(\theta^{n-1}, \frac{\partial L_{cf}}{\partial \theta}, \eta)$ 
8:      $t_u^n \leftarrow \text{StochasticGradientStep}(t_u^{n-1}, \frac{\partial L_{cf}}{\partial t_u}, \lambda_t, \eta)$ 
9:   else
10:     $((u, i, j)) \leftarrow \text{RandomCandidatePair}(P)$ 
11:     $L_{rk} = \ell(s_{u,i} \succ s_{u,j}; \theta)$ 
12:     $\theta^n \leftarrow \text{StochasticGradientStep}(\theta^{n-1}, \frac{\partial L_{rk}}{\partial \theta}, \eta)$ 
13:   end if
14: end for

```

3.3 Implementation of the Joint Approach

We can see that the proposed joint approach is a general framework which can be applied to any existing learning to rank methods. In the following, we adopt two state-of-the-art ranking-based recommendation methods, i.e., BPRMF and HRM, and extend them to the corresponding dynamic-K versions, namely DK-BPRMF and DK-HRM, respectively.

DK-BPRMF Bayesian personalized ranking matrix factorization (BPRMF) is a matrix factorization model with BPR[10] as optimization criteria. With matrix factorization, the target matrix X is approximated by the matrix product of two low-rank matrices $P : |U| \times f$ and $Q : |I| \times f$:

$$\hat{X} := PQ^t \quad (7)$$

where f is the dimensionality/rank of the approximation. Thus the prediction formula can also be written as:

$$\hat{x}_{ui} = \langle p_u, q_i \rangle = \sum_{f=1}^F p_{uf} \cdot q_{if} \quad (8)$$

BPR is a state-of-the-art pairwise ranking framework for the implicit feedback data, specifically, it uses the maximum posterior estimator that is derived from a Bayesian analysis of the problem. Therefore, the ranking objective can be written as:

$$\ell_{rk} = \sum_{(u,i,j) \in P} \sigma(x_{ui} - x_{uj}) + \lambda_\theta \|\theta\|^2 \quad (9)$$

where σ is the logistic sigmoid $\sigma(x) := \frac{1}{1+e^{-x}}$, and λ_θ is a regularization parameter that controls the complexity of the model. According to our joint approach,

the classification objective with logistic loss can be written as :

$$\ell_{cf} = \sum_{(u,i,y_{ui}) \in D} \ln(1 + e^{-y_{ui}(x_{ui}-t_u)}) + \lambda_t \|t_u - t\|^2 \quad (10)$$

Therefore, the dynamic-K version of BPRMF is:

$$\ell_{DK-BPRMF} = \alpha \cdot \ell_{cf} + (1 - \alpha) \cdot \ell_{rk} \quad (11)$$

DK-HRM HRM is a state-of-the-art model for next basket recommendation which can capture both sequential behavior and users' general taste, and meanwhile model complicated interactions among these factors in prediction.

For each user u , a purchase history T^u is given by $T^u := (T_1^u, T_2^u, \dots, T_{t_u-1}^u)$, where $T_t^u \subseteq I$, $t \in [1, t_u - 1]$. Given this history, the task of next basket recommendation is to recommend items that user u would probably buy at the next visit. We define the set H as the set of quaternion $\{(u, t, i, j)\}$ selected from dataset as follows: $H = \{(u, t, i, j) | i \in T_t^u \wedge j \notin T_t^u\}$. The set D is a set of quaternion $\{(u, t, i, y_{u,t,i})\}$ selected from $U \times T \times I$, and $y_{u,t,i}$ is defined as follows:

$$y_{u,t,i} = \begin{cases} 1 & i \in T_t^u \\ -1 & i \notin T_t^u \end{cases} \quad (12)$$

Given a user u and his/her two consecutive transactions T_{t-1}^u and T_t^u , HRM defines the score of buying next item i as following:

$$x_{u,t,i} = v_i \cdot v_{u,t-1}^{Hybrid} \quad (13)$$

where v_i denotes the representation of item i and $v_{u,t-1}^{Hybrid}$ denotes the hybrid representation obtained from the hierarchical aggregation according to T_{t-1}^u .

In the learning process, the ranking objective with logistic loss is written as follows:

$$\ell_{rk} = \sum_{(u,t,i,j) \in H} \frac{1}{1 + e^{-(x_{u,t,i} - x_{u,t,j})}} \quad (14)$$

The classification objective is as follows:

$$\ell_{cf} = \sum_{(u,t,i,y_{u,t,i}) \in D} \ln(1 + e^{-y_{u,t,i}(x_{u,t,i}-t_u)}) + \lambda_t \|t_u - t\|^2 \quad (15)$$

Finally, the dynamic-K version of HRM is:

$$\ell_{DK-HRM} = \alpha \cdot \ell_{cf} + (1 - \alpha) \cdot \ell_{rk} \quad (16)$$

4 Evaluation

In this section, we conduct empirical experiments to demonstrate the effectiveness of our proposed joint learning models DK-BPRMF and DK-HRM for

dynamic-K recommendation. We first introduce the dataset, baseline methods, and the evaluation metrics employed in our experiments. Then we compare the performance of hyper parameters for dynamic-K recommendation. After that, we compare our DK-BPRMF and DK-HRM to the corresponding fixed-N versions to demonstrate effectiveness.

4.1 Datasets

We use two datasets. The Movielens100K dataset and a retail datasets Ta-Feng.

- The Movielens100K dataset, it contains 100,000 ratings (1-5) from 943 users on 1682 movies. Each user has rated at least 20 movies.
- The Ta-Feng dataset is a public dataset released by RecSys conference, which covers products from food, office supplies to furniture. It contains 817, 741 trans- actions belonging to 32, 266 users and 23, 812 items.

We first conduct some pre-process on these datasets. For Movielens100K, as we want to solve an implicit feedback task, we removed the rating scores from the dataset. Now the task is to predict if a user is likely to rate a movie. For Ta-Feng dataset, we remove all the items bought by less than 10 users and users that has bought in total less than 10 items. In order to simulate the real world recommendation scenatio, finally, we seprate the training data and testing data by time. For Ta-Feng, the testing set contains only the last transaction of each user, and for Movielens100K, the testing set contains only the last rated movies of each user, while all the remaining interactions are put into the training set. The models are then learned on S_{train} and their predicted personalized ranking is evaluated on the test set S_{test} .

4.2 Evaluation Methodology

The performance is evaluated for each user u in the testing set. For each recommendation method, we generate a list of K items for each user u , denoted by $R(u)$, where $R_i(u)$ stands for the item recommended in the i -th position. We use the following measures to evaluate the recommendation lists.

- F1-score: F1-score is the harmonic mean of precision and recall, which is a widely used measure in recommendation:

$$Precision(R(u), S_{test}^u) := \frac{|R(u) \cap S_{test}^u|}{|R(u)|} \quad (17)$$

$$Recall(R(u), S_{test}^u) := \frac{|R(u) \cap S_{test}^u|}{|S_{test}^u|} \quad (18)$$

$$F1-score := \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (19)$$

- NDCG@k: Normalized Discounted Cumulative Gain (NDCG) is a ranking based measure which takes into account the order of recommended items in the list[11], and is formally given by:

$$NDCG@k = \frac{1}{N_k} \sum_{j=1}^k \frac{2^{I(R_j(u) \in T_{t_u}^u)} - 1}{\log_2(j+1)} \quad (20)$$

where $I(\cdot)$ is an indicator function and N_k is a constant which denotes the maximum value of NDCG@k given $R(u)$.

- Cover-Ratio: In dynamic-K recommendation, Cover-Ratio measures how many users can get recommendation results.

$$Cover-Ratio = \frac{\sum I(|R(u)| > 0)}{|U|} \quad (21)$$

4.3 Baseline and Parameters Setting

The baselines include

- LogMF[5]: LogMF is a probabilistic model for matrix factorization with implicit feedback.
- BPRMF[10]: BPRMF is a matrix factorization model with BPR as optimize criteria. BPR is a state-of-the-art pairwise ranking framework for the implicit feedback data.
- CRRMF[12]: CRRMF is a matrix factorization model with CRR as optimize criteria. CRR is a hybrid framework combining ranking and regression together.
- HRM[16]: HRM is a state-of-the-art hybrid model on next basket recommendation. Both sequential behavior and users' general taste are taken into account for prediction and meanwhile modeling complicated interactions among these factors in prediction.

For all the models, we fixed the dimension of latent factors in U and I to be 50. In LogMF, similar to Johnson[5], for each triple $(u, i, j) \in P$, i represents positive observations, j represents negative observations and we tune the parameter c to balance the positive and negative observations. In CRRMF, for each triple $(u, i, j) \in P$, we regress i and j to 1 and 0 respectively. In HRM and DK-HRM, we use the average pooling. Parameter tuning was performed using cross validation on the training data for each method. We update parameters utilizing Stochastic Gradient Ddescent(SGD) until converge.

4.4 Comparison among Different Hyper Parameters

For the proposed method, we have three hyper parameters, t , λ_t and α . Because the limitation of space, we only report the result of DK-HRM over two datasets Ta-Feng and Movielens100K. Similar results can also be obtained from DK-BPRMF. The results over two datasets are shown in Figure 1.




Fig. 1: Performance comparison among hyper parameters of DK-HRM over two datasets.

Firstly, we explored how the parameter t affects the performance of DK-HRM. By setting λ_t and α to be 1.0 and 0.5 respectively, we selected t from $\{0.5, 1.0, 1.5, 2.0, 2.5\}$. As we can see, as t increases, F1-score increases but Cover-ratio descreses, indicating that the higher global constraint of personalized decision boundary we set, the higher F1-score we get, but the lower Cover-ratio we get. Take Ta-Feng as an example, when compared with $t = 0.5$, the relative improvement of F1-score by $t = 2.0$ is around 30%, and the decrease of Cover-ratio is 35%.

Then, we fixed t as 2.0 and compared different value of λ_t from $\{0.01, 0.1, 1, 10, 100\}$. As we can see, with the increase of λ_t , which means the increase of regularization strength, F1-score gradually increases, reaches the peak and then starts to decrease. It indicates that appropriate regularization strength can improve metric performance. Take Ta-Feng as example, F1-score reaches peak when $\lambda_t = 1.0$, then starts to decrease.

At last, we fixed t as 2.0 and λ_t as 1.0 and selected α from $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. α denotes the trade-off between optimizing ranking loss and classification loss. As we can see, as α increases, F1 score and NDCG increase first, and then decrease, indicating that the two types of losses can complement each other and achieve the best performance at the same time.

4.5 Comparison against Baselines

We further compare our DK-BPRMF and DK-HRM methods to the state-of-the-art baseline methods. Since LogMF, BPRMF, CRRMF and HRM get top-N recommendation results, in order to make a fair comparison between top-N and dynamic-K, we vary the number N and get the F1-score of top-N method from top-1 to top-20, shown in Figure 2. We find that as the number of recommendation items increases, F1-score increases and then keeps steady. Take Ta-Feng as example, for HRM, F1-score increases to 5.2% and then keeps steady.




Fig. 2: The top-N F1-score of the Ta-Feng dataset and MovieLens dataset. The number of recommended items is set 1 to 20 on Ta-Feng and MovieLens.

Table 1: Performance comparison of DK-HRM and DK-BPRMF among LogMF, BPRMF, CRRMF and HRM over two datasets. Significant improvement of our model(DK-BPRMF and DK-HRM) with respect to their original baseline methods(BPR and HRM respectively)is indicated as ‘+’(p-value ≤ 0.05).

Models	Ta-Feng			MovieLens		
	F1-score	Cover-ratio	NDCG	F1-score	Cover-ratio	NDCG
BPRMF	0.033	1.0	0.075	0.018	1.0	0.067
LogMF	0.042	1.0	0.071	0.028	1.0	0.058
CRRMF	0.041	1.0	0.074	0.026	1.0	0.068
DK-BPRMF	0.048⁺	0.69	0.11⁺	0.021⁺	0.9	0.135⁺
HRM	0.052	1.0	0.08	0.03	1.0	0.117
DK-HRM	0.06⁺	0.7	0.12⁺	0.035⁺	0.68	0.12⁺

Then, we choose the best performed DK-BPRMF(t is set 1.0, α is set 0.5 and λ_t is set 1.0) and DK-HRM(t is set 2.0, α is set 0.5 and λ_t is set 1.0) and the best performed baselines from Figure 2 as the representative for fair comparison. Results over Ta-Feng and MovieLens100K are shown in Table 1. Three main observations can be drawn from the results. (1) BPRMF performs better in NDCG and LogMF performs better in F1-score respectively, due to their different optimization goal for ranking and classification respectively. As a hybrid model, CRRMF can achieve both good ranking and classification performance compared with BPRMF and LogMF. (2) HRM outperforms BPRMF, LogMF and CRRMF in all metrics since it can capture both sequential behavior and users’ general taste. Take Ta-Feng as example, F1-score and NDCG of HRM reach 5.2% and 0.08 respectively. (3) Compared with top-N recommenders, DK-BPRMF and DK-HRM for dynamic-K recommendation can consistently outperform their original baseline methods respectively(i.e. BPR and HRM) in terms of F1-score and NDCG over the two datasets. Take the Ta-Feng dataset as an example, when compared with the second best performed baseline method (i.e. HRM), the relative improvement by DK-HRM(t set as 2.0, α set as 0.5 and λ_t set as 1.0) is around 33.1%, 50%, in terms of F1-score and NDCG, respectively.

It indicates that our joint learning approach for dynamic-K recommendation is more effective than top-N recommendation methods.

5 Conclusions

In this paper, firstly, we argue that it is critical to provide a dynamic-K recommendation where the K should be different with respect to the candidate item set and the target user instead of fixed-N recommendation. Then we formulate this dynamic-K recommendation task as a joint learning problem with both ranking and personalized classification objectives. We apply the above joint learning idea over two state-of-the-art ranking-based recommendation methods, i.e., BPRMF and HRM, and extend them to the corresponding dynamic-K versions, namely DK-BPRMF and DK-HRM, respectively. Our experimental results show that the dynamic-K models are more effective than the corresponding original fixed-N recommendation methods.

References

1. Adomavicius, G., Tuzhilin, A.: Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering* 17(6), 734–749 (2005)
2. Aiolli, F.: Convex auc optimization for top-n recommendation with implicit feedback. In: *Proceedings of the 8th ACM Conference on Recommender Systems*. pp. 293–296. ACM (2014)
3. Feipeng, Z., Yuhong, G.: Improving top-n recommendation with heterogeneous loss. *Journal of Artificial Intelligence Research* (2016)
4. Gopalan, P., Hofman, J.M., Blei, D.M.: Scalable recommendation with poisson factorization. *arXiv preprint arXiv:1311.1704* (2013)
5. Johnson, C.C.: Logistic matrix factorization for implicit feedback data. *Advances in Neural Information Processing Systems* 27 (2014)
6. Karatzoglou, A., Baltrunas, L., Shi, Y.: Learning to rank for recommender systems. In: *Proceedings of the 7th ACM conference on Recommender systems*. pp. 493–494. ACM (2013)
7. Mnih, A., Teh, Y.W.: Learning item trees for probabilistic modelling of implicit feedback. *arXiv preprint arXiv:1109.5894* (2011)
8. Oard, D.W., Kim, J., et al.: Implicit feedback for recommender systems. In: *Proceedings of the AAAI workshop on recommender systems*. pp. 81–83 (1998)
9. Park, D., Neeman, J., Zhang, J., Sanghavi, S., Dhillon, I.S.: Preference completion: Large-scale collaborative ranking from pairwise comparisons. *stat* 1050, 16 (2015)
10. Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: Bpr: Bayesian personalized ranking from implicit feedback. In: *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*. pp. 452–461. AUAI Press (2009)
11. Ricci, F., Rokach, L., Shapira, B.: *Introduction to recommender systems handbook*. Springer (2011)
12. Sculley, D.: Combined regression and ranking. In: *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 979–988. ACM (2010)
13. Shalev-Shwartz, S., Singer, Y., Srebro, N.: Pegasos: Primal estimated sub-gradient solver for svm. In: *Proceedings of the 24th international conference on Machine learning*. pp. 807–814. ACM (2007)
14. Shi, Y., Karatzoglou, A., Baltrunas, L., Larson, M., Hanjalic, A., Oliver, N.: Tfmap: optimizing map for top-n context-aware recommendation. In: *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. pp. 155–164. ACM (2012)
15. Shi, Y., Karatzoglou, A., Baltrunas, L., Larson, M., Oliver, N., Hanjalic, A.: Climf: learning to maximize reciprocal rank with collaborative less-is-more filtering. In: *Proceedings of the sixth ACM conference on Recommender systems*. pp. 139–146. ACM (2012)
16. Wang, P., Guo, J., Lan, Y., Xu, J., Wan, S., Cheng, X.: Learning hierarchical representation model for nextbasket recommendation. In: *Proceedings of the 38th International ACM SIGIR conference on Research and Development in Information Retrieval*. pp. 403–412. ACM (2015)
17. Weimer, M., Karatzoglou, A., Le, Q.V., Smola, A.: Maximum margin matrix factorization for collaborative ranking. *Advances in neural information processing systems* pp. 1–8 (2007)
18. Yun, H., Raman, P., Vishwanathan, S.: Ranking via robust binary classification. In: *Advances in Neural Information Processing Systems*. pp. 2582–2590 (2014)