

Item Group Based Pairwise Preference Learning for Personalized Ranking

Shuang Qiu, Jian Cheng, Ting Yuan, Cong Leng, Hanqing Lu
National Laboratory of Pattern Recognition
Institute of Automation, Chinese Academy of Sciences
{shuang.qiu, jcheng, tyuan, cong.leng, luhq}@nlpr.ia.ac.cn

ABSTRACT

Collaborative filtering with implicit feedbacks has been steadily receiving more attention, since the abundant implicit feedbacks are more easily collected while explicit feedbacks are not necessarily always available. Several recent work address this problem well utilizing pairwise ranking method with a fundamental assumption that a user prefers items with positive feedbacks to the items without observed feedbacks, which also implies that the items without observed feedbacks are treated equally without distinction. However, users have their own preference on different items with different degrees which can be modeled into a ranking relationship. In this paper, we exploit this prior information of a user's preference from the nearest neighbor set by the neighbors' implicit feedbacks, which can split items into different item groups with specific ranking relationships. We propose a novel PRIGP(Personalized Ranking with Item Group based Pairwise preference learning) algorithm to integrate item based pairwise preference and item group based pairwise preference into the same framework. Experimental results on three real-world datasets demonstrate the proposed method outperforms the competitive baselines on several ranking-oriented evaluation metrics.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Information Filtering

Keywords

Collaborative filtering; Implicit feedback; Pairwise preference; Item group

1. INTRODUCTION

Recommender systems, as an increasingly critical tools in dealing with information overload on the Internet, have attracted immense amounts of research recently. Collaborative

Filtering(CF) [5, 7, 9], which is a content-free recommendation technique, widely performs as a core in recommender systems. The underlying assumption of CF is that users with common interests in the past would behave much more similarly on items in the future. Most recent literature of CF is focused on improving the accuracy of regression on users' ratings [3, 9] or personalized ranking of items [10] by effectively exploiting explicit feedbacks (e.g., typically numerical ratings).

However, in most real-world scenarios, explicit feedbacks are not necessarily always available while the abundant implicit or one-class feedbacks [5, 7] are more easily collected such as "clicks" on web pages or "bought" in e-commercial websites. Some pointwise methods [2, 5] regard the observed implicit feedback on an item as a user rating 1 and the unobserved as 0 such that the problem can be addressed by utilizing CF algorithms for explicit feedbacks with various weighting or sampling strategies. Recently, some pairwise methods [7, 12] are proposed with a more proper assumption that a user prefers items with observed feedbacks to items without feedbacks. Note that the assumption of pairwise preference essentially attempts to cope with the issue of personalized ranking directly, which leads to better recommendation performance than pointwise methods. One of the most widely studied pairwise methods is bayesian personalized ranking(BPR) [7] which shows a promising result in handling problems of recommendation with implicit feedbacks. And with its success, various extended methods from distinct aspects for BPR are proposed, e.g., it is extended from two dimensions to three dimensions in [8] by adopting the concept of tensors and scenarios of multiple domains are handled in [4] applying collective matrix factorization.

However, there still exists a challenging problem, which lies in the assumption of pairwise preference over two items. Previous pairwise methods (e.g., BPR) are simply built on the assumption that a user prefers an item with a positive feedback to an item without an observed feedback, which also implies that items without observed feedbacks are treated with no distinction or preference. Nevertheless, a user would have his/her own preference on different items though the items were not observed by him/her before. As a matter of fact, it is reasonable to believe that a user's preference on different items is usually presented in specific ranking relations, which indicates that a user has individual prior preference on different items with different degrees. A recommender system should tend to rank higher those items with greater prior preference for a user. Thereby, prior information for a user should be exploited for a better recommendation.

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 ACM 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.
SIGIR '14, July 6–11, 2014, Gold Coast, Queensland, Australia.
Copyright 2014 ACM 978-1-4503-2257-7/14/07 ...\$15.00.

In this paper, we propose a novel approach which incorporates users' prior preference, which were not considered before, into the algorithms based on the pairwise preference assumption of item level and construct a unified framework. In a typical CF system, a user's nearest neighbors (most similar users according to historical feedbacks) usually share similar interests with that user. Then the items with more positive implicit feedbacks from a user's nearest neighbor set could be more possibly liked by him/her, such that a user's prior preference on items can be discovered. The core concept *item group* is defined in our work as a subset of items with the same cumulative number of implicit feedbacks from nearest neighbor set. Thus, ranking relations of items for a user from his/her prior preference can be reflected by the item groups with different positive feedback numbers. Our approach constructs an item group based pairwise preference for the specific ranking relations of items and combine it with item based pairwise preference to formalise a novel framework PRIGP (Personalized Ranking with Item Group based Pairwise preference learning). Experiments on three real-world datasets demonstrate the effectiveness of our model.

2. OUR APPROACH

2.1 Problem Statement

In a real world scenario, we have a set of users \mathcal{U} with size n and a set of items \mathcal{I} with size m . Each user $u \in \mathcal{U}$ has expressed their positive feedbacks on a set of items $\mathcal{I}_u^+ \subseteq \mathcal{I}$. Moreover, in a user based system, a user u can have fairly similar behaviors with a set of users who constitute the nearest neighbor set denoted as \mathcal{N}_u with size of $|\mathcal{N}_u|$. Consequently, based on the implicit feedbacks of users and information from neighbors, our goal is to recommend a personalized ranking list of items to a user from the item set, $\mathcal{I} \setminus \mathcal{I}_u^+$, that have no observed feedbacks.

2.2 Item Based Pairwise Preference

The basic assumption of pairwise preference of two items can be formally represented as

$$\hat{r}_{ui} > \hat{r}_{uj}, i \in \mathcal{I}_u^+, j \in \mathcal{I} \setminus \mathcal{I}_u^+$$

where \hat{r}_{ui} denotes preference of a user u on an item i and the relation $\hat{r}_{ui} > \hat{r}_{uj}$ indicates user u could have more interests on the item i with a positive feedback than on the item j without an observed feedback. By this item-level assumption, we can further infer the ordinal relations for items in $\mathcal{I} \setminus \mathcal{I}_u^+$ and recommend top- K ranked items to user u .

There are various methods [7, 12] based on the assumption of pairwise preference over two items, which are developed from different perspectives of classification [7] or regression [12]. In our framework, to model the item based pairwise preference, we adopt a similar formulation employed in BPR [7]. Note that our approach is a framework that is not limited to applying BPR.

BPR attempts to maximize the joint probabilities of users' preference on items in their corresponding \mathcal{I}_u^+ more than in $\mathcal{I} \setminus \mathcal{I}_u^+$. By derivations in BPR [7] with the negative log likelihood loss, we can eventually formalize the item based pairwise preference as minimizing the following criterion

$$\text{BPR-Opt} = - \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}_u^+} \sum_{j \in \mathcal{I} \setminus \mathcal{I}_u^+} \ln \sigma(\hat{x}_{uij}(\Theta)) + \mathcal{R}(\Theta) \quad (1)$$

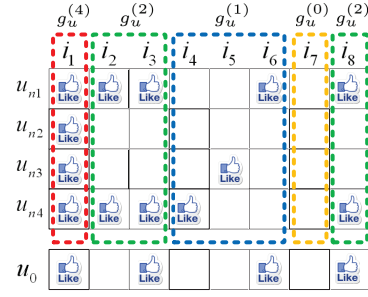


Figure 1: Illustration of Item Group Based Pairwise Preference. The set $\{u_{n1}, \dots, u_{n4}\}$ are the nearest neighbor set of u_0 . Items are split into $g_u^{(0)}$, $g_u^{(1)}$, $g_u^{(2)}$ and $g_u^{(4)}$ while $g_u^{(3)} = \emptyset$, according to implicit feedbacks from neighbors. Thus we have the relationships $\hat{r}_{u_{g_u^{(0)}}} < \hat{r}_{u_{g_u^{(1)}}} < \hat{r}_{u_{g_u^{(2)}}} < \hat{r}_{u_{g_u^{(4)}}}$.

where $\hat{x}_{uij}(\Theta)$ is computed as $\hat{r}_{ui}(\Theta) - \hat{r}_{uj}(\Theta)$ with Θ to be parameters for representing \hat{r}_{ui} and $\sigma(x) = \frac{1}{1+e^{-x}}$ is the logistic sigmoid function defining the likelihood of pairwise preference. More generally, by denoting $\mathcal{L}(\mathcal{I}_u^+, \mathcal{I} \setminus \mathcal{I}_u^+) = - \sum_{i \in \mathcal{I}_u^+} \sum_{j \in \mathcal{I} \setminus \mathcal{I}_u^+} \ln(\sigma(\hat{x}_{uij}(\Theta)))$, we can further formulate the item based pairwise preference learning as

$$\min_{\Theta} \sum_{u \in \mathcal{U}} \mathcal{L}(\mathcal{I}_u^+, \mathcal{I} \setminus \mathcal{I}_u^+) + \mathcal{R}(\Theta) \quad (2)$$

where $\mathcal{L}(\mathcal{I}_u^+, \mathcal{I} \setminus \mathcal{I}_u^+)$ is regarded as a loss function that measures the item based pairwise ranking loss for user u and $\mathcal{R}(\Theta)$ a regularization term to prevent overfitting.

2.3 Item Group Based Pairwise Preference

Previous item based pairwise methods improperly assume that items in the unobserved subset are treated without differences. However, a user may show his/her own preference on different sets of items with different degrees, though the items were not observed by him/her before. Since the recommender system will provide a ranking list of the unobserved items to the users for personalized recommendation, it should tend to rank these items higher for that user. For each user, there can exist a nearest neighbor set that consists of users who behave similarly with that user. Therefore, effectively exploiting the collaborative information from the nearest neighbor set can serve to discover users' prior preference information. In this section, we will show how to infer this prior preference information of a specific user from nearest neighbor set and propose a framework to make personalized ranking by incorporating item group based pairwise preference and the item based pairwise preference.

We define *item group* of a user u as the set of items that have the same number of observed feedbacks from his/her nearest neighbors, which can be formulated as follows

$$g_u^{(k)} = \{i | \mathcal{O}_u(i) = k, i \in \mathcal{I}, k \in \mathbb{Z}\}, \quad (3)$$

where $g_u^{(k)}$ is u 's item group with k feedbacks for each member in it, $\mathcal{O}_u(i)$ denotes the number of feedbacks from u 's nearest neighbor set, \mathbb{Z} is the set of integers and k varies from 0 to $|\mathcal{N}_u|$. Thereby, we assume that a user u is likely to express more preference on an item group with more observed feedbacks from neighbors than an item group with less. We denote $\mathcal{G}_u = \{g_u^{(0)}, g_u^{(1)}, \dots, g_u^{(|\mathcal{N}_u|)}\}$ where \mathcal{G}_u is the set of item groups for user u . Then, the preference relations

among item groups w.r.t. a user $u \in \mathcal{U}$ can be formalised as

$$\hat{r}_{ug_u^{(0)}} < \hat{r}_{ug_u^{(1)}} < \dots < \hat{r}_{ug_u^{|\mathcal{N}_u|}} \quad (4)$$

where $\hat{r}_{ug_u^{(k)}}$ denote the preference of u on his/her k^{th} item group. If there exists a $g_u^{(k)}$ whose size is 0, we can just ignore that item group in Eq. 4. For the above method, we can approximate a user’s prior preference with different degree on different item groups via implicit feedbacks of nearest neighbors. The item group based pairwise preference can be explained by a toy example illustrated in Fig. 1. Specifically, nearest neighbors are determined by cosine similarity of two users’ implicit feedbacks.

Then, we will show how to model item group based pairwise preference of a user. For a user $u \in \mathcal{U}$, we can define u ’s preference on the k^{th} item group as $\hat{r}_{ug_u^{(k)}} = \sum_{i \in g_u^{(k)}} \hat{r}_{ui} / |g_u^{(k)}|$ where $|g_u^{(k)}|$ is the size of $g_u^{(k)}$. Then, we can formulate Eq. 4 into a likelihood probability form of pairwise preference over item groups for each user as

$$P(\mathcal{G}_u) = \prod_{s,t:t>s} p(\hat{r}_{ug_u^{(s)}} > \hat{r}_{ug_u^{(t)}}) \quad (5)$$

where $p(\hat{r}_{ug_u^{(s)}} > \hat{r}_{ug_u^{(t)}})$ is the likelihood probability between item groups $g_u^{(s)}$ and $g_u^{(t)}$. By introducing the negative log likelihood loss function, we construct an item group based pairwise ranking loss in the following form

$$\begin{aligned} \mathcal{L}(\mathcal{G}_u) &= -\ln P(\mathcal{G}_u) \\ &= -\sum_{s,t:t>s} \ln \sigma(\hat{x}_{ug_u^{(t)}g_u^{(s)}}(\Theta)) \end{aligned} \quad (6)$$

where $\hat{x}_{ug_u^{(t)}g_u^{(s)}}(\Theta) = \hat{r}_{ug_u^{(t)}}(\Theta) - \hat{r}_{ug_u^{(s)}}(\Theta)$.

Combining the item based and item group based pairwise ranking loss, we can eventually get our unified framework as

$$\min_{\Theta} \sum_{u \in \mathcal{U}} [\mathcal{L}(\mathcal{I}_u^+, \mathcal{I} \setminus \mathcal{I}_u^+) + \alpha \mathcal{L}(\mathcal{G}_u)] + R(\Theta) \quad (7)$$

where α is a tradeoff parameter used to control the confidence of the pairwise preference among item groups w.r.t. a user and $R(\Theta)$ is a regularization term. Therefore, our model PRIGP are obtained. Note that our model is a general framework and it will reduce to the original item based pairwise methods if $\alpha = 0$ or $|\mathcal{N}_u| = 0$.

More specifically, in our model, the parameter $\Theta = \{U_u \in \mathbb{R}^{1 \times d}, V_i \in \mathbb{R}^{1 \times d}, b_i \in \mathbb{R}, u \in \mathcal{U}, i \in \mathcal{I}\}$ where U_u, V_i are the user and item latent factor vectors respectively and b_i is the item bias. And the parameter regularization term is computed as $R(\Theta) = \frac{1}{2} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} [\lambda_u \|U_u\|^2 + \lambda_v \|V_i\|^2 + \mu_v \|b_i\|^2]$. Thereby, the preference on the item i of a user u is generated by $\hat{r}_{ui} = U_u \cdot V_i^T + b_i$. Moreover, we can compute u ’s preference on the k^{th} item group as $\hat{r}_{ug_u^{(k)}} = U_u \cdot \bar{V}_{g_u^{(k)}}^T + \bar{b}_{g_u^{(k)}}$ where $\bar{V}_{g_u^{(k)}} = (\sum_{i \in g_u^{(k)}} V_i / |g_u^{(k)}|)$ and $\bar{b}_{g_u^{(k)}} = \sum_{i \in g_u^{(k)}} b_i / |g_u^{(k)}|$. Once the parameter Θ are learned, we can predict the preference scores of user u on item $i \in \mathcal{I} \setminus \mathcal{I}_u^+$ as $\hat{r}_{ui} = U_u \cdot V_i^T + b_i$, and a personalized ranking list of items can be recommended to u by sorting predicted scores.

For obtaining a better optimization and boosting the training speed, bootstrapping based stochastic gradient descent method is introduced to learn our model in Eq. 7. We omit the detailed iteration algorithm due to the space limit, and the similar updating method is described in [6, 7].

3. EXPERIMENTS

3.1 Datasets and Baselines

To empirically evaluate our method on the recommendation with implicit feedbacks, we perform experiments on three real-world datasets: MovieLens100K¹, Douban² and Ciao³. MovieLens100K is a widely used benchmark dataset for CF. Douban is a well-known website for users to express their preference on movies, books and music, where we crawled users’ feedbacks on movies. A subset of published Ciao dataset is also collected. To simulate the implicit feedbacks, we keep the ratings larger than 3 as observed positive feedbacks [11] and then we obtain 55,375 observations from 943 users and 1,682 items in MovieLens100K, 276,619 observations from 4,184 users and 2,069 items in Douban and 43,966 observations from 1,296 users and 3,008 items in Ciao.

Two popular baseline methods are used for empirical comparison, which is PopRank [6] and BPR [7]. (1)PopRank is a basic algorithm for the problem of CF with implicit feedbacks, which makes the recommendation to users in terms of global popularity of items. (2)BPR is a strong pairwise method for our comparison and is also a particular case of our method. For the pointwise methods [5, 2], they perform much worse than the pairwise methods (e.g., BPR) and comparisons are not made with them.

3.2 Metrics and Experimental Setup

Considering the real-world scenarios where new users are usually only willing to check a few top-ranked recommended items, to measure the recommendation performance, we adopt the widely used evaluation metrics in top- K recommendation: *NDCG*, *1-call* [1], *F1-score*, *Precision*, *Recall* and with setting $K = 5$. Higher values on above metrics correspond to better recommendation performance.

In our experiments, we randomly select 50% observations as training data and the rest 50% as test data. The dimension of the latent factor vector is $d = 10$. The regularization parameters of our model are set as $\lambda_u = \lambda_v = \mu_v = 0.02, 0.02, 0.1$ for MovieLens100K, Douban and Ciao respectively. And we vary $\alpha \in \{0.01, 0.1, 1, 10, 30\}$ to look for the optimal tradeoff parameter. In addition, we set the size of the nearest neighbours $|\mathcal{N}_u|$ from 0 to 200 with a step 20 to study its impact. We choose the best parameters based on the *NDCG@5* performance on the test set.

3.3 Results and Analysis

Firstly, we compare the performance of our PRIGP approach with other baseline methods as illustrated in Table 1. The optimal parameters we obtain for PRIGP are $\{k = 1, |\mathcal{N}_u| = 140\}$ for MovieLens100K, $\{k = 10, |\mathcal{N}_u| = 200\}$ for for Douban, $\{k = 1, |\mathcal{N}_u| = 40\}$ for Ciao, and detailed discussion will presented in this paper. We observe that our proposed PRIGP framework consistently outperforms all the other baselines. Note that PopRank strategy has a worse performance than BPR and PRIGP and, as a matter of fact, PopRank is not a personalized algorithm for recommendation, which shows the necessity of designing an appropriate personalized recommendation model. A trend

¹<http://grouplens.org/datasets/movielens/>

²<http://movie.douban.com/>

³<http://www.public.asu.edu/~jtang20/datasetcode/>

Table 1: Prediction performance(mean \pm std.) of PopRank, BPR and PRIGP on MovieLens100K, Douban and Ciao datasets. By setting fixing $d = 10$ and $K = 5$. The results in bold indicate the best ones.

Dataset	Method	$NDCG@5$	$1-call@5$	$Prec@5$	$Rec@5$	$F1@5$
MovieLens100K	PopRank	0.3007 \pm 0.0021	0.6777 \pm 0.0196	0.2833 \pm 0.0006	0.0581 \pm 0.0038	0.0872 \pm 0.0041
	BPR	0.4289 \pm 0.0073	0.8558 \pm 0.0120	0.4180 \pm 0.0074	0.1043 \pm 0.0029	0.1478 \pm 0.0037
	PRIGP	0.4452 \pm 0.0013	0.8647 \pm 0.0083	0.4283 \pm 0.0013	0.1076 \pm 0.0016	0.1522 \pm 0.0019
Douban	PopRank	0.3672 \pm 0.0014	0.8028 \pm 0.0020	0.3450 \pm 0.0010	0.0597 \pm 0.0003	0.0984 \pm 0.0005
	BPR	0.4152 \pm 0.0016	0.8430 \pm 0.0043	0.3981 \pm 0.0011	0.0689 \pm 0.0002	0.1135 \pm 0.0002
	PRIGP	0.4389 \pm 0.0031	0.8700 \pm 0.0042	0.4188 \pm 0.0015	0.0724 \pm 0.0002	0.1195 \pm 0.0003
Ciao	PopRank	0.1259 \pm 0.0003	0.3862 \pm 0.0027	0.1236 \pm 0.0007	0.0400 \pm 0.0008	0.0578 \pm 0.0010
	BPR	0.1482 \pm 0.0010	0.4498 \pm 0.0054	0.1414 \pm 0.0012	0.0449 \pm 0.0005	0.0650 \pm 0.0008
	PRIGP	0.1533 \pm 0.0015	0.4730 \pm 0.0076	0.1471 \pm 0.0021	0.0470 \pm 0.0007	0.0680 \pm 0.0009




Figure 2: Parameter Study of $|\mathcal{N}_u|$ and α . (a)(b)(c) study the $|\mathcal{N}_u|$ on MovieLens100K, Douban and Ciao while (d)(e)(f) study α on MovieLens100K, Douban and Ciao.

of our PRIGP methods can be observed that it always outperforms BPR, which proves the effectiveness of our methods for personalized recommendation issues with implicit feedbacks. Moreover, since our method adopts BPR into its framework, the comparison of PRIGP and BPR could demonstrate that by discovering the prior information of a user’s preference from nearest neighbor set, our method can definitely improve the pairwise personalized ranking method the recommendation performance.

Then, the impact of parameter $|\mathcal{N}_u|$ is studied explicitly with the metric $NDCG@5$. As illustrated in Fig. 2(a),(b),(c), by fixing optimal α for all datasets, we can observe that the recommendation performance will firstly increase rapidly as the size gets larger because more information from neighbors can be obtained. Then with $|\mathcal{N}_u|$ becoming larger, the curves in Fig. 2(a),(c) reach an optimal point and turn to decrease, and the curve in Fig. 2(b) increases much slower. That is because a larger nearest neighbor set will include more users with a lower similarity such that the information from neighbors will not increase or even more noisy than helpful ranking information will be involved from that neighbor set. Specifically, the optimal $|\mathcal{N}_u|$ for MovieLens100K, Douban and Ciao are at 140, 200, 40 respectively.

Finally, to understand the effect of different nearest neighbor set sizes, we study the influence of the tradeoff parameter α which balances the item based and item group based pairwise preference in ranking. As illustrated in Fig. 2(d), (e), (f), we can observe that $\alpha = 1, 10, 1$ are the most proper tradeoff parameters empirically with setting $|\mathcal{N}_u|$ to be the best one in each dataset.

4. CONCLUSIONS

In this paper, we propose a novel recommendation algorithm PRIGP for CF problems with implicit feedbacks. In our model, we integrate item based pairwise preference and item group based pairwise preference into the same framework with leveraging the observed feedbacks from neighbors. The experimental results on three real-world datasets show our proposed method performs a better top- K recommendation than baseline methods. Noting that our work provides a framework which can be fit for any personalized ranking method, we plan to generalize it to other pairwise methods in the future.

5. ACKNOWLEDGMENTS

This work was supported in part by 973 Program (Grant No. 2010CB327905), National Natural Science Foundation of China (Grant No. 61170127, 61332016).

6. REFERENCES

- [1] H. Chen and D. R. Karger. Less is more: probabilistic models for retrieving fewer relevant documents. In *SIGIR*, 2006.
- [2] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM*, 2008.
- [3] Y. Koren. Factor in the neighbors: Scalable and accurate collaborative filtering. *ACM TKDD*, 4(1), 2010.
- [4] A. Krohn-Grimberghe, L. Drumond, C. Freudenthaler, and L. Schmidt-Thieme. Multi-relational matrix factorization using bayesian personalized ranking for social network data. In *WSDM*, 2012.
- [5] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *ICDM*, 2008.
- [6] W. Pan and L. Chen. Cofiset: Collaborative filtering via learning pairwise preferences over item-sets. *SDM*, 2013.
- [7] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, 2009.
- [8] S. Rendle and L. Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *WSDM*, 2010.
- [9] R. Salakhutdinov and A. Mnih. Probabilistic matrix factorization. In *NIPS*, 2007.
- [10] Y. Shi, M. Larson, and A. Hanjalic. List-wise learning to rank with matrix factorization for collaborative filtering. In *RecSys*, 2010.
- [11] V. Sindhwani, S. Bucak, J. Hu, and A. Mojsilovic. A family of non-negative matrix factorizations for one-class collaborative filtering problems. In *RecSys*, 2009.
- [12] G. Takács and D. Tikk. Alternating least squares for personalized ranking. In *RecSys*, 2012.