TRecSo: Enhancing Top-k Recommendation With Social Information *

Chanyoung Park, Donghyun Kim, Jinoh Oh, Hwanjo Yu†
Dept. of Computer Science and Engineering
POSTECH (Pohang University of Science and Technology), South Korea
{pcy1302, kdh5377, kurin, hwanjoyu}@postech.ac.kr

ABSTRACT

Due to the data sparsity problem, social network information is often additionally used to improve the performance of recommender system. While most existing works exploit social information to reduce the rating prediction error, e.g., RMSE, a few had aimed to improve the top-k ranking prediction accuracy. This paper proposes a novel top-k oriented recommendation method, TRecSo, which incorporates social information into recommendation by modeling two different roles of users as trusters and trustees while considering the structural information of the network. Empirical studies on real-world datasets demonstrate that TRecSo leads to remarkable improvement compared to previous methods in top-k recommendation.

Keywords

Recommender System; Social network; Learning-To-Rank

1. INTRODUCTION

Recommending top-k items is the eventual goal in typical recommender system rather than accurately predicting the ratings of all items, as users are only interested to see top-k items [1]. Most recommender systems, however, mainly focus on accurately predicting the overall ratings [6, 5, 2, 12], and they are not well optimized for the task of finding top-k items. Several methods have been developed based on the Learning-To-Rank (LTR) perspective to provide accurate results at top-k [10, 14]. However, they still suffer from the data sparsity problem, that is, the recommendation is hardly accurate due to lack of observations (i.e., ratings) because users typically rate a small number of items. To tackle the data sparsity problem, researchers have tried to incorporate auxiliary information such as social network relationship, text reviews on items, etc. This paper focuses on incorporating the social network information of users in the top-k recommendation.

Two top-k recommendation methods have been developed to incorporate social network information based on the LTR approach [13, 14]. Specifically, Yao *et al.* [13] linearly combines a user's taste and her direct friends' tastes in optimizing the top-k recommendation. However, it does not utilize other important information hidden in social network such as the structural information or truster-trustee relationship. Zhao *et al.* [14] optimizes the top-k recommendation from relative ordering that can be extracted from purchase history or browsing history, but it cannot handle numerical ratings directly. Note that numerical ratings usually contain much richer information on user's preference than relative ordering.

This paper proposes a novel LTR-based top-k recommendation method, TRecSo, which leverages the social network information to optimize top-k recommendation. TRecSo is distinguished from previous methods in that it models two different roles of users as trusters and trustees while considering the structural information of the network. Our experimental results on real-world datasets indicate that TRecSo considerably outperforms previous methods in top-k recommendation. Our implementation and experiment results can be found in our technical report [9].

2. METHOD

We first explain, in Section 2.1, how the social information is incorporated in TRecSo, and describe the objective function to optimize top-k recommendation in Section 2.2.

2.1 Incorporating Social Information

Assume that there are N users and M items, and $R = [r_{ij}]$ is a $N \times M$ matrix where r_{ij} represents the rating that user i gave on item j. The rating matrix R is typically very sparse whose entries are mostly unknown. Then, the rating of user i on item j is predicted as follows:

$$\hat{r}_{ij} = g(\mu + b_{U_i} + b_{q_j} + q_j{}^T(\alpha p_i + (1-\alpha)w_i + |I_i|^{-\frac{1}{2}}\sum_{t\in I_i} y_t + |T_i|^{-\frac{1}{2}}\sum_{v\in T_i} x_v)) \quad (1)$$

where $g(\cdot)$ is the logistic function that bounds the range of predicted ratings; μ is the average of all ratings; b_{U_i} and b_{q_j} represent the user and item biases, respectively; q_j represents the item latent vector; p_i and w_i represent the user latent vectors as truster and as trustee, respectively, which are also used to model the social information in Eq.(2); α balances between the truster and trustee roles; I_i denotes the set of items rated by user i; y_t is the latent vector of item t, which models implicit influence of items rated by user i; T_i is the set of users that user i trusts (e.g., whom user i follows in social network); and x_v is the latent vector of whom user i

*This work was supported by the Industrial Core Technology Development Program (10049079 , Development of Mining core technology exploiting personal big data) funded by the Ministry of Trade, Industry and Energy (MOTIE, Korea), ICT R&D program of MSIP/IITP [14-824-09-014, Basic Software Research in Human-level Lifelong Machine Learning (Machine Learning Center)]
†Corresponding author

WWW'16 Companion, April 11–15, 2016, Montréal, Québec, Canada.
ACM 978-1-4503-4144-8/16/04.
http://dx.doi.org/10.1145/2872518.2889362 .

Table 1: The NDCG@5 and NDCG@10 averaged over 5 runs. The best performance is in bold.

Dataset	FilmTrust				Ciao				Epinion			
	N=20		N=50		N=20		N=50		N=20		N=50	
NDCG	5	10	5	10	5	10	5	10	5	10	5	10
Itemknn	0.580	0.639	0.564	0.636	0.694	0.740	0.667	0.714	0.623	0.676	0.599	0.651
OCCF[8]	0.618	0.673	0.620	0.669	0.708	0.755	0.681	0.722	0.641	0.691	0.635	0.673
BPRMF[10]	0.620	0.680	0.575	0.641	0.705	0.749	0.646	0.693	0.625	0.681	0.619	0.659
ListRank[11]	0.673	0.718	0.618	0.652	0.682	0.737	0.661	0.702	0.692	0.739	0.635	0.673
SBPR[14]	0.632	0.683	0.570	0.629	0.702	0.752	0.647	0.696	0.628	0.680	0.617	0.660
SoRank[13]	0.667	0.714	0.642	0.675	0.679	0.736	0.680	0.719	0.656	0.706	0.654	0.692
TRecSo	**0.684**	**0.729**	**0.653**	**0.700**	**0.785**	**0.818**	**0.753**	**0.790**	**0.738**	**0.776**	**0.678**	**0.717**

trusts, which models implicit influence of the users trusted by user i. $\hat{r}_{ij}$ in Eq.(1) will be used to optimize top-k recommendation in Section 2.2.

To tackle the data sparsity problem, social network information is modeled in TRecSo as follows. Given an asymmetric social relation matrix $S = \{s_{iv}\}, [s_{iv}] \in \{0, 1\}$ where s_{iv} indicates whether user i trusts (or follows) user v or not, the unknown relationship $\hat{s}_{iv}$ between user i and v can be estimated as follows:

$$\hat{s}_{iv} = g\left(b_{p_i} + b_{w_v} + {w_v}^T p_i\right) \tag{2}$$

where b_{p_i} and b_{w_v} represent the truster bias and trustee bias, respectively. By sharing the term p_i and w_v in Eq.(1) and (2), and by *simultaneously learning both latent models*, we are able to properly model the different roles of users as trusters and trustees. Note that our method can be generalized to the case where the social relations are symmetric. i.e., Friendship.

To reflect the structural information of the network, s_{iv} is adjusted based on the degree of nodes such that it gives lower weights to those who *trust* many users and gives higher weights to those who *are trusted* by many users:

$$s^*_{iv} = \sqrt{\frac{Indegree(v_v)}{Outdegree(v_i) + Indegree(v_v)}} \times s_{iv} \tag{3}$$

where v_i and v_v are nodes for user i and user v in the network, respectively [5]. It is worth noting that the model performance has actually improved by this adjustment in the experiments [9].

2.2 Top-k Optimization

To optimize top-k recommendation, we formulate our objective based on top-one probability, $P_{l_i}(C_{ij}) = \frac{exp(C_{ij})}{\sum_{k=1}^{K} exp(C_{ik})}$, which models the probability of an item scored C_{ij} being ranked on the top-one position in user i's ranked list l_i [1]. By utilizing the top-one probability, we are now able to formulate the objective function aiming at minimizing the uncertainty between the training list and the predicted list by using cross-entropy measure as follows, which can be interpreted as list-wise ranking prediction:

$$\begin{aligned} L = &-\sum_i \sum_{j \in I_i} P_{li}(r_{ij}) \log P_{li}(\hat{r}_{ij}) - \lambda_t \sum_i \sum_{v \in T_i} P_{li}(s^*_{iv}) \log P_{li}(\hat{s}_{ij}) \\ &+ \frac{\lambda_b}{2}(||b_{U_i}||_F^2 + ||b_{q_j}||_F^2 + ||b_{p_i}||_F^2 + ||b_{w_v}||_F^2) \\ &+ \frac{\lambda}{2}(||p_i||_F^2 + ||w_i||_F^2 + ||q_j||_F^2 + \sum_i ||y_i||_F^2 + \sum_v ||x_v||_F^2) \end{aligned} \tag{4}$$

where λ_t is a parameter that controls the importance of social regularization. Having formulated the non-convex objective function as shown in Eq.(4), we compute the gradient of each latent vector, i.e., $p_i, q_j, y_t, w_v, x_v, b_{U_i}, b_{p_i}, b_{q_j}, b_{w_v}$, and learn them by stochastic gradient descent [7] from which we obtain the local minimum.

3. EXPERIMENTS

Dataset: We used three public real-world datasets for evaluation (Table 2). The social relations between users are asymmetric in all three datasets. More details about data statistics and experiment results can be found in our technical report [9].

Table 2: Data Statistics

Dataset	Users	Items	Ratings	Density	Trust
FilmTrust	1,508	2,071	35,497	1.1400%	1,853
Ciao	7,375	99,746	278,483	0.0379%	111,781
Epinion	40,163	139,738	664,824	0.0118%	487,183

Setup: We compared TRecSo with six state-of-the-art methods that fall into one of three categories: 1) **Traditional CF**: *ItemKnn*, 2) **Ratings-only-based methods**: *OCCF*, *BPRMF*, *ListRank*, and 3) **Social network based methods**: *SBPR*, *SoRank*. We set $\alpha = 0.5$, $\lambda = 0.01, \lambda_b = 0.01$ and $\lambda_t = 0.8$ for TRecSo, and the parameters for all the other baselines are set to their best performing parameters. Note that latent dimensions of 5 and learning rate of 0.01 are used for all the experiments. For fair comparison, we used the same experimental protocol as in [13].

Results: Table 1 shows that our method, TRecSo, consistently outperforms all the state-of-the-art methods for all datasets. Note that N in Table 1 denotes the number of items for each user in the training data.

4. CONCLUSION

This paper proposes TRecSo, a novel matrix factorization based recommendation method that optimizes the top-k ranking prediction accuracy by additionally considering the social network information. Specifically, TRecSo integrates the social network information into the Learning-To-Rank (LTR) based objective function for recommendation. Comprehensive experimental results show that TRecSo significantly outperforms the state-of-the-art algorithms in the top-k ranking accuracy of recommendation.

5. REFERENCES

[1] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In *ICML*, 2007.

[2] P. Cremonesi, Y. Koren, and R. Turrin. Performance of recommender algorithms on top-n recommendation tasks. In *RecSys*. ACM, 2010.

[3] G. Guo, J. Zhang, D. Thalmann, and N. Yorke-Smith. Etaf: An extended trust antecedents framework for trust prediction. In *ASONAM*, 2014.

[4] G. Guo, J. Zhang, and N. Yorke-Smith. A novel bayesian similarity measure for recommender systems. In *IJCAI*, 2013.

[5] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *CIKM*. ACM, 2008.

[6] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *NIPS*, pages 1257–1264, 2007.

[7] J. Oh, W.-S. Han, H. Yu, and X. Jiang. Fast and robust parallel sgd matrix factorization. In *KDD*. ACM, 2015.

[8] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *ICDM*, 2008.

[9] C. Park, D. Kim, J. Oh, and H. Yu. TRecSo: Enhancing top-k recommendation with social information. In *http://dm.postech.ac.kr/TRecSo*, 2016.

[10] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, 2009.

[11] Y. Shi, M. Larson, and A. Hanjalic. List-wise learning to rank with matrix factorization for collaborative filtering. In *RecSys*. ACM, 2010.

[12] W. Yao, J. He, G. Huang, and Y. Zhang. Modeling dual role preferences for trust-aware recommendation. In *SIGIR*. ACM, 2014.

[13] W. Yao, J. He, G. Huang, and Y. Zhang. Sorank: incorporating social information into learning to rank models for recommendation. In *WWW*, 2014.

[14] T. Zhao, J. McAuley, and I. King. Leveraging social connections to improve personalized ranking for collaborative filtering. In *CIKM*, 2014.