Exploiting Latent Social Listening Representations for Music Recommendations

Chih-Ming Chen[1,2], Po-Chuan Chien[1], Yu-Ching Lin[3], Ming-Feng Tsai[1], Yi-Hsuan Yang[2]

[1]Department of Computer Science
National Chengchi University
Taipei 11605, Taiwan
{99303052, mftsai}@nccu.edu.tw

[2]Research Center for Information Technology Innovation
Academia Sinica
Taipei 11564, Taiwan
{cmchen, yang}@citi.sinica.edu.tw

[3]Research Center
KKBOX Inc.
Taipei 11501, Taiwan
aaronlin@kkbox.com

ABSTRACT

Music listening can be regarded as a social activity, in which people can listen together and make friends with one other. Therefore, social relationships may imply multiple facets of the users, such as their listening behaviors and tastes. In this light, it is considered that social relationships hold abundant valuable information that can be utilized for music recommendation. However, utilizing the information for recommendation could be difficult, because such information is usually sparse. To address this issue, we propose to learn the latent social listening representations by the DeepWalk method, and then integrate the learned representations into Factorization Machines to construct better recommendation models. With the DeepWalk method, user social relationships can be transformed from the sparse and independent and identically distributed (*i.i.d.*) form into a dense and non-*i.i.d.* form. In addition, the latent representations can also capture the spatial locality among users and items, therefore benefiting the constructed recommendation models.

Keywords

Representation Learning, Factorization Machine, Recommender System, Social Network, Graph

1. BACKGROUND

Underlying almost all recommendation algorithms is an attempt to model the interaction among users and items. There have been some studies working on utilizing auxiliary information for improving recommendations. In [2, 5], social relationships are utilized to densify the ratings of users to improve the similarity computation behind the Collaborative Filtering-based (CF-based) methods. For the Matrix Factorization-based (MF-based) methods, some studies focus on how to incorporate social relations with other attributes of users [3, 4, 10] and how to affect the regularization term [6, 11]. In addition, the Collaborative Topic Regression (CTR) methods [1, 8] fuse the idea of topic modeling with the probabilistic matrix factorization on social networks, in order to infer useful latent topics for collaborative filtering.

2. METHODOLOGY

Figure 1 illustrates the framework of the proposed method. We transfer the listening history and friends relationship into a social listening graph. Given the social listening graph, DeepWalk is used to learn the implicit representation by seeking the possible path on the graph. The idea is to maximize the co-occurrence patterns in each generated path so that the potential distance is modeled. The representation, usually presented as a vector, can encode the context information for further processing [7]. From theoretical perspective, the social representation learning is a technique of combining the recent developments of language modeling and unsupervised representation learning. In this work, we use the technique to learn the representations on a social listening graph, and then feed the learned representations into Factorization Machines (FM) [9]. After being processed by DeepWalk, the sparse social relations will become dense and in a non-*i.i.d.* form, which can be helpful for increasing the connections among users and items. To our best knowledge, this work is the first attempt to use deepwalk for music recommendation.

To find the possible path, we utilize random walk to uniformly sample a series of random vertex from a graph. The primitive graph is constructed on the sole social network. For the specific application of music recommendation, we propose 3 different ways to construct such a social graph:

- **Social** graph: Build the graph only based on users' friends, which will enable DeepWalk to detect local community.
- **Listening** graph: Build the graph only based on user-item listening matrix, which will enable DeepWalk to identify the association patterns about users and items.
- **Social Listening** graph: Build the graph based on the above two relations including the user-user and user-item matrices, which may hopefully fuse the merits of the preceding two approaches.

3. EXPERIMENTS

Our experiments involve two real-world music datasets – hetrec2011-lastfm-2k and KKBOX-50K. The first one is a public benchmark dataset derived Last.fm. The second one is collected from a music streaming company KKBOX. Table 2 shows some statistics of the datasets.

We randomly hold out 80% records for each user as the training data. The remaining 20% records of all users are

RecSys 2015 Poster Proceedings, September 16-20, 2015, Austria, Vienna.

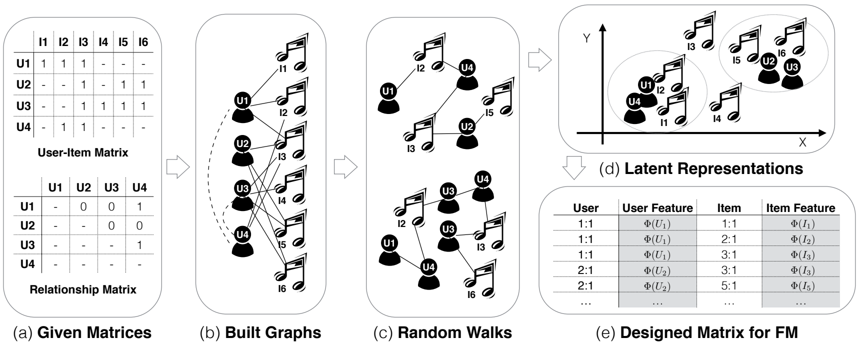

Figure 1: Exploring latent social listening representations for music recommendation.

		hetrec2011-lastfm-2k			KKBOX-50K		
Approach	DeepWalk	MAP	Recall@100	Recall@200	MAP	Recall@100	Recall@200
Sole User-to-Item	✗	4.993%	29.554%	40.308%	5.167%	4.498%	8.143%
Friendship Indexes	✗	5.204%	32.275%	39.111%	2.764%	2.031%	3.971%
Random Social Graph	✓	4.993%	30.479%	41.089%	0.311%	0.004%	0.004%
Social Graph	✓	5.646%	33.328%	45.153%	5.155%	4.494%	8.038%
Listening Graph	✓	8.423%	46.285%	60.486%	6.094%	5.376%	9.658%
Social Listening Graph	✓	**8.807%**	**47.962%**	**62.113%**	**6.157%**	**5.446%**	**9.708%**

Table 1: Recommendation Performance

Dataset	#User	#Item	#Record	#Relations
lastfm-2k	1,892	17,632	92,834	12,717
KKBOX-50K	50,000	200,000	50,000,000	80,000

Table 2: Data Statistics

treated as the testing item pool. Instead of using pure precision, we use recall and mean average precision (MAP) as the performance measurements. We repeat the evaluation process 5 times with different randomly selected training sets and report the average performance.

Since FM model is already a promising approach of MF-based model, we directly perform all the experiments based on FM, including a) one primitive approach that uses pure user-to-item matrix, b) one baseline approach that directly uses friendship indexes as auxiliary feature c) one testing approach that builds the graph by randomly generated social links and d) three proposed approaches that integrate different types of latent representation in correspondingly. Evaluation results are reported in Table 1.

Social links are much more sparser than the case of user-to-item ratings so that direct use of friendship is ineffective. Besides, combing the social graph with listening graph is able to filter out the inadequate connections. Consequently, we can see that the use of social listening graph leads to about 60% relative improvement to baseline approach (*i.e.* Friendship Indexes) for hetrec2011-lastfm-2k dataset and about 20% of improvement for KKBOX-50K dataset in terms of MAP and recall measurements. A randomly generated social graph is also examined for verifying the effectiveness of social relations in music data. It can be found that the random graph leads to inferior result than the proposed ones for building the graph, integrating the latent representation learned from the social listening graph achieves the best improvement in both MAP and recall measurements.

4. CONCLUSIONS AND FUTURE WORK

We propose a novel method that bridges unsupervised representation learning for social links and context-aware factorization model for recommendations. According to the experimental results, direct use of friendship indices as features does not perform well. Among the three proposed ways

In current work, all the connections are considered to have the same weight (*i.e.* the binary response) and some types of connections are omitted (*i.e.* the item-to-item connections). Hence, it is possible to earn better performance by, for example, assigning numeric weights to the connections or adding item-to-item connections. We leave these as future work.

5. REFERENCES

[1] C. Chen, X. Zheng, Y. Wang, F. Hong, and Z. Lin. Context-aware collaborative topic regression with social matrix factorization for recommender systems. In *Proc. AAAI*, 2014.

[2] W.-S. Hwang, S. Li, S.-W. Kim, and K. Lee. Data imputation using a trust network for recommendation. In *Proc. WWW Companion*, 2014.

[3] M. Jiang, P. Cui, R. Liu, Q. Yang, F. Wang, W. Zhu, and S. Yang. Social contextual recommendation. In *Proc. ACM CIKM*, 2012.

[4] X. Liu and K. Aberer. Soco: A social network aided context-aware recommender system. In *Proc. WWW*, 2013.

[5] C. Luo, W. Pang, Z. Wang, and C. Lin. Hete-cf: Social-based collaborative filtering recommendation using heterogeneous relations. In *Proc. IEEE ICDM*, 2014.

[6] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *Proc. ACM WSDM*, 2011.

[7] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proc. ACM KDD*, 2014.

[8] S. Purushotham, Y. Liu, and C. J. Kuo. Collaborative topic regression with social matrix factorization for recommendation systems. *CoRR*, 2012.

[9] S. Rendle. Factorization machines with libFM. *ACM Trans. Intell. Syst. Technol.*, 2012.

[10] X. Yang, H. Steck, and Y. Liu. Circle-based recommendation in online social networks. In *Proc ACM KDD*, 2012.

[11] Q. Yuan, L. Chen, and S. Zhao. Factorization vs. regularization: Fusing heterogeneous social relationships in top-n recommendation. In *Proc. ACM RecSys*, 2011.