

Decentralized Recommender Systems

Zhangyang Wang* Xianming Liu* Shiyu Chang* Jiayu Zhou† Guo-Jun Qi‡
 Thomas S. Huang*

Abstract

This paper proposes a *decentralized recommender system* by formulating the popular collaborative filtering (CF) model into a decentralized matrix completion form over a set of users. In such a way, data storages and computations are fully distributed. Each user could exchange limited information with its local neighborhood, and thus it avoids the centralized fusion. Advantages of the proposed system include a protection on user privacy, as well as better scalability and robustness. We compare our proposed algorithm with several state-of-the-art algorithms on the FlickrUserFavor dataset, and demonstrate that the decentralized algorithm can gain a competitive performance to others.

1 Introduction

The paper discusses the decentralized recommender systems, which is in contrast to the typical recommender systems built on centralized infrastructures (the “cloud”, etc.). The decentralized network [1] had been thoroughly investigated in control and communication fields, defined as a set of distributed but connected agents, who are generally not strongly connected in a graph theoretic sense. Each agent collects data by itself, and executes computation via limited communication within only local neighborhoods. Fig. 1 shows a comparison of centralized versus decentralized network structures. Specifically, in a decentralized recommender system, individual users / user-groups can be viewed as network agents. Each user holds his or her own ratings as partially observed data. The data cannot be accessed by any intermediate point or centralized server in the network. Therefore, it has a potential effect on protecting user data privacy against both the cloud server and some malicious eavesdropping over uploading channels. For a large-scale network of mobile users, the decentralized models own a better scalability since users are only locally connected. Since data storages and computations are fully distributed, the decentralized systems also become robust to center (cloud) or individual agent (user) failures.

To our best knowledge, we are the first studying and designing a decentralized recommender system. There have




Figure 1: The comparison of a centralized network (left) and a decentralized network (right).

been some recent interests in investigating algorithms in a decentralized fashion [2, 3], and moreover, preliminary literatures to reveal the convergence [4] and dynamics [5] properties. All above make it solid and promising to build application scenes on a decentralized network structure.

2 Model and Algorithm

As a most popular tool in recommender system, Collaborative Filtering (CF) is usually formulated as matrix factorizations problem [7] [11]. It predicts user rating $\mathbf{R}_{i,j}$ (of i -th user on j -th item) as a dot product of the user profile of the i -th user, denoted as a row vector \mathbf{U}^i , and the item profile of the j -th item denoted as a column vector \mathbf{V}_j , i.e., $\mathbf{R}_{i,j} = \mathbf{U}^i \mathbf{V}_j$. The recommendation problem can be formulated as solving the following matrix factorization problem:

$$(2.1) \quad \begin{aligned} \min_{\mathbf{U}, \mathbf{V}, \mathbf{Z}} \frac{1}{2} \|\mathbf{UV} - \mathbf{Z}\|_2^2 \\ \text{s.t. } P_{\Omega}(\mathbf{Z}) = P_{\Omega}(\mathbf{R}) \end{aligned}$$

Here P_{Ω} denotes the projection over the set of available ratings, and \mathbf{Z} is an auxiliary matrix.

It is assumed that CF is performed by L users jointly in a decentralized manner, and \mathbf{R} is segmented into L non-overlapped parts, denoted as \mathbf{R}_i , $i = 1, 2, \dots, L$. For example, the easiest case to segment \mathbf{R} is to divide by columns. The i -th user ($i = 1, 2, \dots, L$) observes \mathbf{R}_i . Note some level of synchronization is still required to collaboratively utilize information from all users. The trade-off strategy is **to share partial data only among users in the local neighborhood**. After observing the problem structure, authors in [2] suggested an variant of nonlinear Gauss-Seidel (GS) iterations, named *decentralized matrix completion (DMC)* algorithm. The i -th user will hold \mathbf{R}_i , as well as \mathbf{U}_i , \mathbf{V}_i , and \mathbf{Z}_i based on its own computations. Note \mathbf{U}_i here is of the same size as \mathbf{U} , and \mathbf{Z}_i of the same dimension as \mathbf{R}_i , so in other words,

*Beckman Institute, University of Illinois at Urbana-Champaign, IL 61821. {zwang119, xliu102, chang87, t-huang1}@illinois.edu.

†Samsung Research America, San Jose, CA 95134. jiayu.zhou@samsung.com

‡University of Central Florida, Orlando, FL, 32816, USA. guojun.qi@ucf.edu.

Algorithm 1 Decentralized matrix completion (DMC) algorithm for solving (2.1)

Require: $P_\Omega(\mathbf{R}_i)$, ($i = 1, 2, \dots, L$); initializations of \mathbf{U}_i , \mathbf{V}_i , and \mathbf{Z}_i for each i -th user ($i = 1, 2, \dots, L$); step size β ; ITER

1: FOR $t=1$ to ITER DO

2: Each i -th user updates \mathbf{V}_i : $\mathbf{V}_i = (\mathbf{U}_i^T \mathbf{U}_i)^{-1} \mathbf{U}_i^T \mathbf{Z}_i$

3: Each i -th user updates \mathbf{Z}_i : $\mathbf{Z}_i = \mathbf{U}_i \mathbf{V}_i + P_\Omega(\mathbf{R}_i - \mathbf{U}_i \mathbf{V}_i)$

4: Each i -th user propagates \mathbf{U}_i to its one-hop neighborhood N_i .

5: Each i -th user updates \mathbf{U}_i :

$$\mathbf{U}_i = \frac{\mathbf{z}_i \mathbf{V}_i^T - \mathbf{a}_i + \beta \sum_{j \in N_i} \mathbf{U}_j}{1 + 2\beta |N_i|}$$

$$\mathbf{a}_i = \mathbf{a}_i + \beta (|N_i| \mathbf{U}_i - \sum_{j \in N_i} \mathbf{U}_j)$$

6: END

Ensure: \mathbf{U}_i , \mathbf{V}_i , and \mathbf{Z}_i , $i = 1, 2, \dots, L$

$\mathbf{Z}_i = \mathbf{U}_i \mathbf{V}_i$. In each iteration, the i -th user first updates \mathbf{V}_i and \mathbf{Z}_i independently, then exchanging \mathbf{U}_i with its one-hop connected neighborhood users, and finally update \mathbf{U}_i via the average consensus algorithm[6]. The algorithm is summarized in Algorithm I. It obtains similar reconstruction errors, compared to centralized solutions [2].

3 Experiments

We compare our proposed algorithm with state-of-the-arts in this section, including Probabilistic Matrix Factorization (PMF) [8], and Collaborative Topic Modeling (CTR) [9], on a collected image recommendation dataset from Flickr.

FlickrUserFavor dataset: The dataset contains 350,000 images collected from Flickr, from 140 user groups and uploaded by 20,298 users. We use the “like” feedback provided by users as binary ratings. 75% of the rating matrix is used as training, and 25% as testing.

Evaluation Measurement: We use the averaged ranked order of all the rated images in the testing dataset for a specific user to evaluate performances. Among these ranked images, we determine those for which a user has exhibited a “like” preference in the *test data*, and report the *average percentile score* (APS) of the ranked images which are indeed preferred by the user. The lower the APS, the better the algorithm is, which means the user preferred images are ranked in top positions. Finally, the *mAPS* is reported with the mean of the APS scores for all target users.

Performance Comparison: PMF [8] is the most classical collaborative filtering algorithm for recommender system. Wang et.al. also proposed the Collaborative Topic Model [9] which involves both content and user ratings, where we use the Hierarchical Gaussianization (HG) [10] as the image features. For the proposed decentralized algorithm, we set the rank as 64, and using 8 agents. Detailed

mAPS comparisons are shown in Table 1. It is shown that DMC is capable to achieve competitive performance as *CTR*, while is far better than *PMF*. Moreover, we do not use any content information in our algorithm, while *CTR* uses content to indicates similarities between items. That suggests a further direction improve our algorithm too.

Table 1: Performances of the proposed approaches compared with other baseline methods. The second column indicates whether or not the algorithm uses content information. Without using content information and the fusion center, the proposed algorithm achieves a competitive performance.

method	Content	mAPS
PMF [8]	N	61.99
CTR [9]	Y	52.76
DMC	N	53.46

4 Conclusion

This paper discusses a decentralized recommender system. We formulate the popular collaborative filtering model into a decentralized matrix completion problem. Each user, with only partial rating data, can exchange user profile factors with its local neighborhood, while keep item profile factors private. We compare our proposed algorithm with several state-of-the-arts on the FlickrUserFavor dataset, and illustrate comparable results to the conventional ones.

5 Acknowledgement

The work is supported by Samsung under the 2013 GRO project “Device centric Social Network Platform”.

References

- [1] C. Yeung, I. Llicardi, K. Lu, O. Seneviratne and T. Lee, “Decentralization: The future of online social networking”, In W3C Workshop on Future of Social Networking Position Papers.
- [2] Q. Ling, Y. Xu, W. Yin and Z. Wen, “Decentralized low-rank matrix completion,” in Proceedings of ICASSP, pp. 2925-2928, 2012.
- [3] Q. Ling, Z. Wen, W. Yin, “Decentralized jointly sparse optimization by reweighted ℓ_q minimization,” IEEE Trans. on Signal Processing, 61.5 (2013): 1165-1170.
- [4] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, “On the linear convergence of the ADMM in decentralized consensus optimization,” (2014): 1-1.
- [5] Q. Ling and A. Ribeiro, “Decentralized dynamic optimization through the alternating direction method of multipliers,” IEEE 14th Workshop on SPAWC, pp. 170-174, 2014.
- [6] T. Erseghe, D. Zennaro, E. Dall’Anese, and L. Vangelista, “Fast consensus by the alternating direction multipliers method,” IEEE Trans. on Signal Processing, 59.11 (2011): 5523-5537.
- [7] Y. Koren, R. Bell, and C. Volinsky, Matrix factorization techniques for recommender systems, Computer, vol. 42, no. 8, pp. 30-37, 2009.
- [8] R. Salakhutdinov and A. Mnih, “Probabilistic matrix factorization,” in Proceedings of NIPS, 2008, vol. 20.
- [9] C. Wang and D. M. Blei, “Collaborative topic modeling for recommending scientific articles,” in Proceedings of SIGKDD, 2011, pp. 448 – 456.
- [10] X. Zhou, N. Cui, Z. Li, F. Liang, and T. S. Huang, “Hierarchical gaussianization for image classification,” in Proceedings of ICCV, pp. 1971–1977, 2009
- [11] S. Chang, G. Qi, C. Aggarwal, J. Zhou, M. Wang, and T. S. Huang, Factorized similarity learning in networks. in Proceedings of ICDM, 2014.