Enhanced Trust SVD: Recommendation Model with Explicit and Implicit Influence of User Trust and Ratings

I.Edwin Albert[1], Dr. Deepa A.J[2]
[1]Research Scholar, Department of Information and Communication Engineering, Anna University, Chennai.

[2]Professor, Department of Electronics and Communication Engineering, Ponjesly College of Engineering, Nagercoil.

Abstract

An Enhanced Trust SVD model, trust-based matrix factorization technique is used for Recommendation system. Enhanced Trust SVD (Singular Value Decomposition) integrates multiple information sources into the recommendation model in order to reduce well known problems Data Sparsity and Cold start Problems and their degradation recommendation performance. By analysing the social trust data from four real-world data sets, we conclude that not only the explicit but also the implicit influence of both ratings and trust should be taken into consideration in a recommendation model. Hence, we build on top of a state-of-the-art recommendation algorithm SVD++ which inherently involves the explicit and implicit influence of rated items, by further incorporating both the explicit and implicit influence of trusted users on the prediction of items for an active user. To our knowledge, the work reported is the first to extend SVD++ with social trust information. Experimental results on the four data sets demonstrate that our approach Enhanced Trust SVD achieves better accuracy than other ten counterparts, and can better handle the concerned issues.

Introduction

Incorporating trust into recommender systems has demonstrated potential to improve recommendation performance (Yang et al. 2016; Fang, Bao, and Zhang 2014), and to help mitigate some well-known issues, such as data sparsity and cold start (Guo, Zhang, and Thalmann 2012). Such trust aware approaches are developed based on the phenomenon that friends often influence each other by recommending items. However, even the best performance reported by the latest work (Fang, Bao, and Zhang 2014) can be inferior to that of other state-of-the-art models which are merely based on user-item ratings. For instance, a well-performing trust-based model (Yang et al. 2016) obtains 1.0585 on data set Epinions.com in terms of Root Mean Square Error (RMSE), whereas the performance of a user-item baseline (see Koren (2017), Sect. 2.1) can achieve 1.0472 in terms of RMSE.1

To investigate this phenomenon, we conduct an empirical trust analysis based on four real-word data sets (FilmTrust, Epinions, Flixster and Ciao) through which two important observations are concluded. First, trust information is also very sparse, yet complementary to rating information. Hence, focusing too much on either one kind of information may achieve only marginal gains in predictive accuracy. Second, users are strongly correlated with their trust neighbours whereas they have a weakly positive correlation with their trust-alike neighbours (e.g., friends). Given that very few trust networks exist, it is better to have a more general trust-based model that can well operate on both trust and trust-alike relationships. These observations motivate us to consider both explicit and implicit influence of ratings and of trust in a trust-based model. The influence can be explicit (real values of ratings and trust) or implicit (who rates what (for ratings) and who trusts whom (for trust)). The implicit influence of ratings has been demonstrated useful in providing accurate recommendations (Koren 2017).We will later show that implicit trust can also provide added value over explicit trust. Thus we propose a novel trust-based recommendation model Trust SVD.

Our approach builds on top of a state-of the- art model SVD++ (Koren 2017) where both the explicit and implicit influence of user-item ratings are involved to generate predictions. To the authors' knowledge, our work is the first to extend SVD++ with social trust information. Specifically, on one hand the implicit influence of trust (who trusts whom) can be naturally added to the SVD++ model by extending the user modelling. On the other hand, the explicit influence of trust (trust values) is used to constrain that user specific vectors should conform to their social trust relationships.

This ensures that user-specific vectors can be learned from their trust information even if a few or no ratings are given. In this way, the data sparsity and cold start issues can be better alleviated. Our novel model thus incorporates both explicit and implicit influence of item ratings as well as user trust. In addition, a weighted regularization technique is used to further avoid over-fitting for model learning. Experimental results on the four real-world data sets demonstrate that our approach achieves significantly better accuracy than other trust-based counterparts as well as other ratings-only well-performing models (ten approaches in total), and is more capable of coping with cold start situations

Related Work

Trust-aware recommender systems have been widely studied, given that social trust provides an alternative view of user preferences other than item ratings (Guo, Zhang, and Yorke-Smith 2014). Specifically, Ma et al. (2017) propose a social regularization method (SoRec) by considering the constraint of social relationships. The idea is to share a common user-feature matrix factorized by ratings and by trust. Ma, King, and Lyu (2018) then propose a social trust ensemble method (RSTE) to linearly combine a basic matrix factorization model and a trust-based neighbourhood model together. Ma et al. (2011) further propose that the active user's user-specific vector should be close to the average of her trusted neighbours, and use it as a regularization to form a new matrix factorization model (SoReg).

Jamali and Ester (2010) build a new model (Social MF) on top of SoRec by reformulating the contributions of trusted users to the formation of the active user's user-specific vector rather than to the predictions of items. Yang et al. (2016) propose a hybrid method (TrustMF) that combines both a trustor model and a trustee model from the perspectives of trustors and trustees, that is, both the users who trust the active user and those who are trusted by the user will influence the user's ratings on unknown items. Tang et al. (2016) consider both global and local trust as the contextual information in their model, where the global trust is computed by a separate algorithm. Yao et al. (2014) take into consideration both the explicit and implicit interactions among trustors and trustees in a recommendation model.

Fang, Bao, and Zhang (2014) stress the importance of multiple aspects of social trust. They decompose trust into four general factors and then integrate them into a matrix factorization model. All these works have shown that a matrix factorization model regularized by trust outperforms the one without trust. That is, trust is helpful in improving predictive accuracy. However, it is also noted that even the latest work (Fang, Bao, and Zhang 2014) can be inferior to other well-performing ratings-only models. To explain this phenomenon, we next conduct a trust analysis to investigate the value of trust in recommender systems.

Trust Analysis : Four data sets are used in our analysis and also our later experiments:
Dataset and its analysis

- The model uses 4 real-world datasets for analysis namely Epinions, FilmTrust, Flixster and Ciao. The trust data is even sparser than the rating data in all datasets except Ciao.

- We observed that even inactive users (in terms of rating) were still socially connected to other users thus helping to solve the cold-start problem. There is a high correlation between a user's rating and the average of her social neighbours.

Observation 1 Trust information is very sparse, yet is complementary to rating information. On one hand, as shown in Table 1, the density of trust is much smaller than that of ratings in Epinions, FilmTrust and Flixster whereas trust is only denser than ratings in Ciao. Both ratings and trust are very sparse across all the data sets.

Table 1: Statistics of the four data sets

Feature	Epinions	FilmTrust	Flixster	Ciao
users	40,163	1,508	53,213	7,375
items	139,738	2,071	18,197	99,746
ratings	664,824	35,497	409,803	280,391
density	0.051%	1.14%	0.04%	0.03%
trusters	33,960	609	47,029	6,792
trustees	49,288	732	47,029	7,297
trusts	487,183	1,853	655,054	111,781
density	0.029%	0.42%	0.03%	0.23%

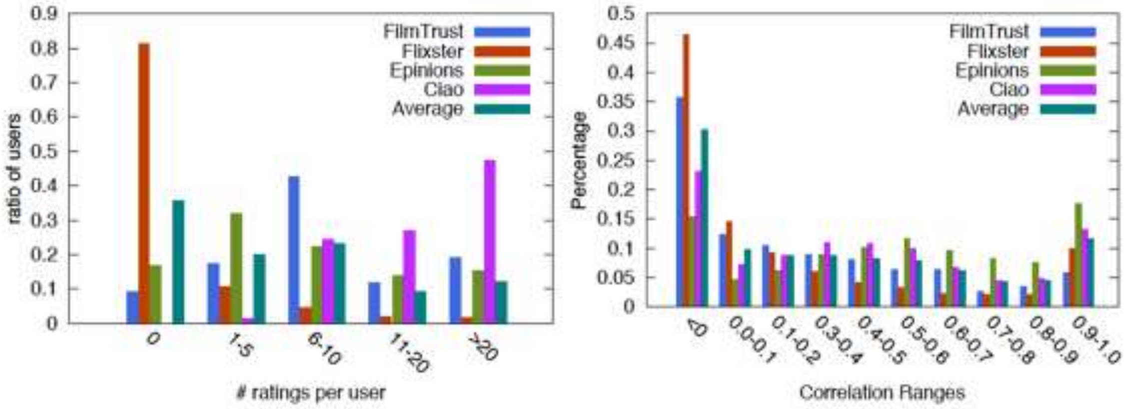

(a) (b)

Figure 1: (a) The distribution of ratio of users who have issued trust statements with respect to the number of ratings that they each have given. (b) The correlations between a user's ratings and those of her trusted neighbours in all the data sets.

Observation 2 A user's ratings have a weakly positive correlation with the average of her social neighbours under the concept of trust-alike relationships, and a strongly positive correlation under the concept of trust relationships. Next, we consider the influence of trust in rating prediction, i.e., the influence of trusted neighbours on the active user's rating for a specific item, a.k.a. social influence. Specifically, we calculate the Pearson correlation coefficient (PCC) between a user's ratings and the average of her social neighbours. The results are presented in Figure 1b, indicating that: (1) A weakly positive correlation is observed between a user's ratings and the average of the social neighbours in FilmTrust (mean 0.183) and Flixster (0.063).

The distributions of the two data sets are similar. Flixster adopts the symmetric friendship relationships whereas trust is directed. Although FilmTrust adopts the concept of trust (with values from 1 to 10), the publicly available data set contains only binary values (such degrading may cause much noise). We regard these relationships as trust alike, i.e., the social relationships that are similar with, but weaker (or more-noisy) than social trust. (2) Under the concept of trust relationships, on the contrary, a user's ratings are strongly and positively correlated with the average of trusted neighbours. Specifically, a large portion (17.63% in Epinions, 13.14% in Ciao) of user correlations are in the range of [0:9; 1:0], and (resp. 54.70%, 39.14%) of user correlations are greater than 0.5. The average correlation is 0.446 in Epinions, and 0.322 in Ciao. Since PCC values are in the range of [-1, 1] values of 0.446 and 0.322 indicate decent correlations. In the social networks with relatively weak trust-alike relationships (e.g., friendship), implicit influence (i.e., binary relationships) may be more indicative than explicit (but noisy) values for recommendations. Hence, a trust-based model that ignores the implicit influence of ratings and trust may lead to deteriorated performance if being applied to such cases.

The second observation suggests that incorporating both the explicit and implicit influence of ratings and trust may promote the generality of a trust-based model to both trust and trust-alike social relationships. Our approach presented next is constructed based on these two observations.

Enhanced Trust SVD: A Trust-based Model

Problem Definition

The reason to define the algorithm for predicting the user's interest instead of existing algorithms are:

a. Collaborative Filtering suffers from two well-known issues are data sparsity and cold start.

b. Unsuitable for real life applications because of the increased computational and communication costs. Some other problems are:

1. Cold start: It's difficult to give recommendations to new users as his/her profile is almost empty and he has not rated any items yet so his taste is unknown to the system. This is called the cold start problem. In some recommender systems this problem is solved with observation when creating a profile. Items may also have a cold-start when they are fresh in the system and haven't been rated before. Both of these problems can be also solved with hybrid approaches.

2. Trust: The voices of people with a short history may not be that relevant as the voices of those who have rich history in their profiles. The issue of trust arises towards evaluations of a definite customer. The issue could be solved by distribution of preferences to the users.

3. Scalability: With the growth of numbers of users and items, the system requires more resources for processing information and forming recommendations. Most of resources is consumed with the purpose of determining users with similar tastes, and goods with similar descriptions. This problem can also be cleared by the combination of several types of filters and physical enhancement of systems. Parts of numerous computations may also be implemented offline in order to accelerate issuance of recommendations online.

4. Sparsity: In online shopping those have a huge number of users and items there are almost always users that have rated just a few items. Using collaborative filtering and other approaches recommender systems generally create neighbourhoods of users using their profiles. If a user has evaluated just few items then it's pretty difficult to determine his/her taste and he/she could be related to the wrong neighbourhood. Sparsity is the problem of lack of information.

5. Privacy: Privacy has been the most important problem. In order to obtain the most accurate and exact recommendation, the system must gain the most amount of information possible about the user, including demographic data, and data about the location of a particular user. Automatically, the question of reliability, security and confidentiality of the given information arises. Many online shops offer effective protection of privacy of the users by utilizing specialized algorithms and programs.

Enhanced Trust SVD

The essence is to find two low-rank matrices (SVD): user-feature matrix (p) (K x M) and item-feature matrix (q) (K x N) which could regenerate R (M x N). K denotes the dimensions of the latent vector space

$$\hat{r}_{u,j} = q_j^\top p_u.$$

Learn the user and item feature matrix my minimizing the following loss function through SGD, ALS, etc. L2 regularisation is applied to the loss function to reduce overfitting.

$$\mathcal{L}_r = \frac{1}{2}\sum_u \sum_{j\in I_u} (q_j^\top p_u - r_{u,j})^2 + \frac{\lambda}{2}\Big(\sum_u \|p_u\|_F^2 + \sum_j \|q_j\|_F^2\Big)$$

SVD++ that including implicit user feedback in the model can lead to a much better prediction. They provided the following prediction rule:

$$\hat{r}_{u,j} = b_u + b_j + \mu + q_j^\top \Big(p_u + |I_u|^{-\frac{1}{2}} \sum_{i\in I_u} y_i\Big)$$

Trust SVD: Incorporates **explicit as well as implicit social trust data** into the SVD++ model.

- We can denote the trust information by the matrix T (M x M) which indicates the extent to which one user trusts another. Then by SVD find two lower rank matrices p (K x M) (trustor feature matrix) and w (K x M) (trustee feature matrix) to learn the trust relationships. Matrix p is shared across R and T (explicit trust learning). Matrices p and w can be learned by minimizing the loss function similar to the equation earlier:

$$\mathcal{L}_t = \frac{1}{2}\sum_{u}\sum_{v \in T_u}\left(w_v^\top p_u - t_{u,v}\right)^2 + \frac{\lambda}{2}\left(\sum_u \|p_u\|_F^2 + \sum_v \|w_v\|_F^2\right)$$

- The final predicted rating with an implicit rating and trust information can be given by the equation below. The additional term represents the user's latent features in terms of the trust relationship with other users. So, an item highly rated by her trusted users would receive higher ratings.

$$\begin{aligned}\hat{r}_{u,j} &= b_u + b_j + \mu \\ &\quad + q_j^\top \left(p_u + |I_u|^{-\frac{1}{2}} \sum_{i \in I_u} y_i + |T_u|^{-\frac{1}{2}} \sum_{v \in T_u} w_v\right)\end{aligned}$$

- The model encourages greater deviations from the baseline estimates for users who provided more ratings or had a better social network and sticks to the baseline predictions for users who have less interaction within the system observed that the model performed better when this behaviour was moderated a little. Hence the terms are divided by the number of elements for a user u in the matrices I and T.

- The final loss function represents a unified recommendation model learning from both types of information. A few additional regularisation terms are added to avoid further overfitting.

$$\begin{aligned}\mathcal{L} = &\tfrac{1}{2}\sum_u \sum_{j \in I_u} \left(\hat{r}_{u,j} - r_{u,j}\right)^2 + \tfrac{\lambda_t}{2}\sum_u \sum_{v \in T_u}\left(\hat{t}_{u,v} - t_{u,v}\right)^2 \\ &+ \tfrac{\lambda}{2}\sum_u |I_u|^{-\frac{1}{2}} b_u^2 + \tfrac{\lambda}{2}\sum_j |U_j|^{-\frac{1}{2}} b_j^2 \\ &+ \sum_u \left(\tfrac{\lambda}{2}|I_u|^{-\frac{1}{2}} + \tfrac{\lambda_t}{2}|T_u|^{-\frac{1}{2}}\right)\|p_u\|_F^2 \\ &+ \tfrac{\lambda}{2}\sum_j |U_j|^{-\frac{1}{2}}\|q_j\|_F^2 + \tfrac{\lambda}{2}\sum_i |U_i|^{-\frac{1}{2}}\|y_i\|_F^2 \\ &+ \tfrac{\lambda}{2}|T_v^+|^{-\frac{1}{2}}\|w_v\|_F^2,\end{aligned}$$

Evaluation metrics

- A five-fold cross-validation is used for training and testing.

- RMSE and MAE are used as error evaluation metrics.

- Two dataset views are created for testing. One which takes in all the ratings (All view) and the other which targets cold start data (particularly users who have rated less than five items) (Cold start view).

Results

- We compare our results with baseline models, state-of-the-art rating only models and some recent models which incorporate trust data as well.

- Betters other models in accuracy except while calculating the MAE metric for Epinions dataset.

All	Metrics	UserAvg	ItemAvg	PMF	RSTE	SoRec	SoReg	SocialMF	TrustMF	SVD++	TrustSVD	Improve
FilmTrust $d = 5$	MAE	0.636	0.725	0.714	0.628	0.628	0.674	0.638	0.631	0.613*	**0.609**	0.65%
	RMSE	0.823	0.927	0.949	0.810	0.810	0.878	0.837	0.810	0.804*	**0.789**	1.87%
	MAE	0.636	0.725	0.735	0.640	0.638	0.668	0.642	0.631	0.611*	**0.607**	0.65%
$d = 10$	RMSE	0.823	0.927	0.968	0.835	0.831	0.875	0.844	0.819	0.802*	**0.787**	1.87%
Epinions $d = 5$	MAE	0.930	0.928	0.979	0.950	0.882	0.994	0.825	0.818	0.818*	**0.804**	1.71%
	RMSE	1.203	1.094	1.290	1.196	1.114	1.315	1.070	1.069	1.057*	**1.043**	1.32%
	MAE	0.930	0.928	0.909	0.958	0.884	0.932	0.826	0.819	0.818*	**0.805**	1.59%
$d = 10$	RMSE	1.203	1.094	1.197	1.278	1.142	1.232	1.082	1.095	1.057*	**1.044**	1.23%
Flixster $d = 5$	MAE	0.729*	0.858	0.814	0.751	0.750	0.820	0.770	0.890	0.794	**0.726**	0.41%
	RMSE	0.979*	1.088	1.076	0.975	0.974	1.087	0.994	1.146	1.062	**0.948**	3.17%
	MAE	0.729*	0.858	0.769	0.784	0.785	0.785	0.784	1.116	0.821	**0.727**	0.27%
$d = 10$	RMSE	0.979*	1.088	1.009	1.015	1.018	1.034	1.009	1.441	1.091	**0.950**	2.97%
Ciao $d = 5$	MAE	0.781	0.760	0.920	0.767	0.765	0.899	0.749	0.742*	0.752	**0.723**	2.56%
	RMSE	1.031	1.026	1.206	1.020	1.013	1.183	0.981*	0.983	1.013	**0.955**	2.65%
	MAE	0.781	0.760	0.822	0.763	0.761	0.815	0.749*	0.753	0.748	**0.723**	3.47%
$d = 10$	RMSE	1.031	1.026	1.078	1.013	1.010	1.076	0.976*	1.014	1.001	**0.956**	2.05%

Cold Start	Metrics	UserAvg	ItemAvg	PMF	RSTE	SoRec	SoReg	SocialMF	TrustMF	SVD++	TrustSVD	Improve
FilmTrust $d = 5$	MAE	0.709	0.722	0.814	0.680	0.670*	0.881	0.697	0.674	0.677	**0.661**	1.34%
	RMSE	0.979	0.911	1.079	0.884	0.857*	1.104	0.916	0.867	0.897	**0.853**	0.47%
	MAE	0.709	0.722	0.767	0.674	0.668*	0.771	0.680	0.687	0.680	**0.663**	0.75%
$d = 10$	RMSE	0.979	0.911	1.009	0.900	0.897*	1.034	0.907	0.900	0.905	**0.853**	4.91%
Epinions $d = 5$	MAE	1.047	0.852*	1.451	1.051	0.892	1.398	0.884	0.891	0.889	**0.868**	-1.88%
	RMSE	1.430	1.127	1.770	1.266	1.138	1.735	1.133	1.125*	1.162	**1.105**	1.78%
	MAE	1.047	0.852	1.153	0.981	0.846*	1.139	0.857	0.853	0.891	**0.868**	-2.60%
$d = 10$	RMSE	1.430	1.127*	1.432	1.313	1.180	1.437	1.152	1.176	1.166	**1.108**	1.69%
Flixster $d = 5$	MAE	0.869	0.906	1.097	0.872	0.872	1.058	0.881	0.901	0.868*	**0.855**	1.50%
	RMSE	1.155	1.114	1.390	1.097	1.096*	1.358	1.103	1.138	1.122	**1.066**	2.74%
	MAE	0.869*	0.906	0.949	0.889	0.892	0.951	0.884	0.976	0.869*	**0.858**	1.27%
$d = 10$	RMSE	1.155	1.114	1.206	1.137	1.144	1.218	1.112*	1.328	1.112*	**1.070**	3.78%
Ciao $d = 5$	MAE	0.829	0.735*	1.033	0.957	0.789	1.173	0.774	0.752	0.759	**0.729**	0.82%
	RMSE	1.138	1.005	1.334	1.113	0.998	1.430	1.001	0.954*	1.039	**0.953**	0.10%
	MAE	0.829	0.735	0.926	0.803	0.730*	0.949	0.741	0.770	0.749	**0.721**	1.23%
$d = 10$	RMSE	1.138	1.005	1.191	1.014	1.031	1.214	0.978*	1.096	1.020	**0.962**	1.64%

- Presents a major improvement in modelling trust info and reinforces the importance of implicit data.
- The model of SVD++ that is used in this paper is the second-best model in terms of performance that is mentioned in the original paper. The best model was the hybrid one which uses the mixture of the latent factor as well the neighbourhood model. Would be interesting to see the results in that case.
- It would be interesting to use this particular data and trust data in general in Factorisation Machines (combines advantages of SVM with Factorisation Models) which specializes in dealing with sparse data.
- Modelling trust for the recommendation task of item ranking instead of rating prediction.

Conclusion and Future Work

We proposed a novel trust-based matrix factorization model which incorporated both rating and trust information. Our analysis of trust in four real-world data sets indicated that trust and ratings are complementary to each other, and both pivotal for more accurate recommendations. Our novel approach, Enhanced Trust SVD, takes into account both the explicit and implicit influence of ratings and trust information when predicting ratings of unknown items. A weighted regularization technique was adapted and used to further regularize the user- and item-specific latent feature vectors. Comprehensive experimental results showed that our approach Enhanced Trust SVD outperformed both trust- and ratings-based methods in predictive accuracy across different testing views and across users with different trust degrees. For future work, we intend to further improve the proposed model by considering both the influence of trustors and trustees.

References

Fang, H.; Bao, Y.; and Zhang, J. 2014. Leveraging decomposed trust in probabilistic matrix factorization for effective recommendation. In Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI), 30–36.

G. Adomavicius and A. Tuzhilin, "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 (TKDE), vol. 17, no. 6, pp. 734–749, 2017.

Guo, G.; Zhang, J.; and Thalmann, D. 2012. A simple but effective method to incorporate trusted neighbours in recommender systems. In Proceedings of the 20th International Conference on User Modelling, Adaptation and Personalization (UMAP), 114–125.

Steffen Rendle: Factorization Machines, in Proceedings of the 10th IEEE International Conference on Data Mining (ICDM 2016), Sydney, Australia.

Guo, G.; Zhang, J.; and Yorke-Smith, N. 2014. Leveraging Multiview's of trust and similarity to enhance clustering based recommender systems. Knowledge-Based Systems (KBS) in press.

Jamali, M., and Ester, M. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In Proceedings of the 4th ACM Conference on Recommender Systems (RecSys), 135–142.

Koren, Y. 2017. Factorization meets the neighbourhood: a multifaceted collaborative filtering model. In Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), 426–434.

Koren, Y. 2010. Factor in the neighbours: Scalable and accurate collaborative filtering. ACM Transactions on Knowledge Discovery from Data (TKDD) 4(1):1:1–1:24.

Ma, H.; Yang, H.; Lyu, M.; and King, I. 2017. SoRec: social recommendation using probabilistic matrix factorization. In Proceedings of the 31st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 931–940.

Ma, H.; Zhou, D.; Liu, C.; Lyu, M. R.; and King, I. 2011. Recommender systems with social regularization. In Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM), 287–296.

Ma, H.; King, I.; and Lyu, M. 2018. Learning to recommend with social trust ensemble. In Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 203–210.

Salakhutdinov, R., and Mnih, A. 2017. Probabilistic matrix factorization. In Advances in Neural Information Processing Systems (NIPS), volume 20, 1257–1264.

Tang, J.; Hu, X.; Gao, H.; and Liu, H. 2016. Exploiting local and global social context for recommendation. In Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), 2712–2718.

Yang, B.; Lei, Y.; Liu, D.; and Liu, J. 2016. Social collaborative filtering by trust. In Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), 2747–2753.

Yao, W.; He, J.; Huang, G.; and Zhang, Y. 2014. Modelling dual role preferences for trust-aware recommendation. In Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR), 975–978.