

A Probabilistic Model for the Cold-Start Problem in Rating Prediction using Click Data

ThaiBinh Nguyen¹ and Atsuhiko Takasu^{1,2}

¹ Department of Informatics,
SOKENDAI (The Graduate University for Advanced Studies), Tokyo, Japan

² National Institute of Informatics, Tokyo, Japan
{binh,takasu}@nii.ac.jp

Abstract. One of the most efficient methods in collaborative filtering is matrix factorization, which finds the latent vector representations of users and items based on the ratings of users to items. However, a matrix factorization based algorithm suffers from the *cold-start* problem: it cannot find latent vectors for items to which previous ratings are not available. This paper utilizes *click data*, which can be collected in abundance, to address the cold-start problem. We propose a *probabilistic item embedding* model that learns item representations from click data, and a model named EMB-MF, that connects it with a *probabilistic matrix factorization* for rating prediction. The experiments on three real-world datasets demonstrate that the proposed model is not only effective in recommending items with no previous ratings, but also outperforms competing methods, especially when the data is very sparse.

Keywords: Recommender system, Collaborative filtering, Item embedding, Matrix factorization

1 Introduction

Rating prediction is one of the key tasks in recommender systems. From research done previously on this problem, matrix factorization (MF) [12, 8, 5] is found to be one of the most efficient techniques. An MF-based algorithm finds the *vector representations* (latent feature vectors) of users and items, and uses these vectors to predict the unseen ratings. However, an MF-based algorithm suffers from the *cold-start* problem: it cannot find the latent feature vectors for items that do not have any prior ratings; thus cannot recommend them.

To address the *cold-start* problem, many methods have been proposed. Most of them rely on exploiting *side information* (e.g., the contents of items). The collaborative topic model [14] and content-based Poisson factorization [6] use text content information of items as side information for recommending new items. In [15], the author proposed a model for music content using deep neural network and combined it with an MF-based model for music recommendation. However, in many cases, such side information is not available, or is not informative enough (e.g., when an item is described only by some keywords).

This paper focuses on utilizing *click data*, another kind of feedback from the user. The advantage of utilizing click data is that it can be easily collected with abundance during the interactions of users with the systems. The idea is to identify item representations from click data and use them for rating prediction.

The main contributions of this work can be summarized as follows:

- We propose a *probabilistic item embedding* model for learning item embedding from click data. We will show that this model is equivalent to performing PMF [12] of the positive-PMI (PPMI) matrix, which can be done efficiently.
- We propose EMB-MF, a model that combines the *probabilistic item embedding* and PMF [12] for coupling the item representations of the two models.
- The proposed model (EMB-MF) can automatically control the contributions of prior ratings and clicks in rating predictions. For items that have few or no prior ratings, the predictions mainly rely on click data. In contrast, for items with many prior ratings, the predictions mainly rely on ratings.

2 Proposed Method

2.1 Notations

The notations used in the proposed model are shown in Table 1.

Table 1. Definitions of some notations

Notation	Meaning
N, M	the number of users, number of items
R, S	the rating matrix, PPMI matrix
\mathcal{R}	the set of (u, i) -pair that rating is observed (i.e., $\mathcal{R} = \{(u, j) R_{uj} > 0\}$)
$\mathcal{R}_u, \mathcal{R}_i$	the set of items that user u rated, set of users that rated item i
\mathcal{S}	the set of (i, j) -pair that $S_{ij} > 0$ (i.e., $\mathcal{S} = \{(i, j) S_{ij} > 0\}$)
\mathcal{S}_i	the set of item j that $S_{ij} > 0$ (i.e., $\mathcal{S}_i = \{j S_{ij} > 0\}$)
d	the dimensionality of the feature space
\mathbf{I}_d	the d -dimensional identity matrix
ρ_i, α_i	the <i>item embedding vector</i> and <i>item context vector</i> of item i
θ_u, β_i	the <i>feature vector</i> of user u , <i>feature vector</i> of item i , in the rating model
μ, b_u, c_i	the <i>global mean</i> of ratings, <i>user bias</i> of user u , <i>item bias</i> of item i
B_{ui}	$b_u + c_i + \mu$
$\boldsymbol{\rho}, \boldsymbol{\alpha}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{b}, \mathbf{c}$	$\{\rho\}_{i=1}^M, \{\alpha\}_{i=1}^M, \{\theta\}_{u=1}^N, \{\beta\}_{i=1}^M, \{b\}_{u=1}^N, \{c\}_{i=1}^M$
$\boldsymbol{\Omega}$	the set of all model parameters (i.e., $\boldsymbol{\Omega} = \{\boldsymbol{\rho}, \boldsymbol{\alpha}, \boldsymbol{\theta}, \boldsymbol{\beta}, \mathbf{b}, \mathbf{c}\}$)

2.2 Probabilistic Item Embedding Based on Click Data

The motivation behind the use of clicks for learning representations of items is the following: if two items are often clicked in the *context* of each other, they are likely to be similar in their nature. Therefore, analyzing the *co-click* information of items can reveal the relationship between items that are often clicked together.

“*Context*” is a modeling choice and can be defined in different ways. For example, the context can be defined as the set of items that are clicked by the user (*user-based* context); or can be defined as the items that are clicked in a session (*session-based* context). Although we use the user-based context to describe the proposed model in this work, other definitions can also be used.

We represent the association between items i and j via a *link function* $g(\cdot)$, which reflects how strong i and j are related, as follows:

$$p(i|j) = g(\rho_i^\top \alpha_j) p(i) \quad (1)$$

where $p(i)$ is the probability that item i is clicked; $p(i|j)$ is the probability that i is clicked by a user given that j has been clicked by that user. We want the value of $g(\cdot)$ to be large if i and j are frequently clicked by the same users.

There are different choices for the link functions, and an appropriate choice is $g(i, j) = \exp\{\rho_i^\top \alpha_j\}$. Eq. 1 can be rewritten as:

$$\log \frac{p(i|j)}{p(i)} = \rho_i^\top \alpha_j \quad (2)$$

Note that $\log \frac{p(i|j)}{p(i)}$ is the point-wise mutual information (PMI) [4] of i and j , and we can rewrite Eq. 2 as:

$$PMI(i, j) = \rho_i^\top \alpha_j \quad (3)$$

Empirically, PMI can be estimated using the actual number of observations:

$$\widehat{PMI}(i, j) = \log \frac{\#(i, j)|\mathcal{D}|}{\#(i)\#(j)} \quad (4)$$

where \mathcal{D} is the set of all item–item pairs that are observed in the click history of all users, $\#(i)$ is the number of users who clicked i , $\#(j)$ is the number of users who clicked j , and $\#(i, j)$ is the number of users who clicked both i and j .

A practical issue arises here: for item pair (i, j) that is not often clicked by the same user, $PMI(i, j)$ is negative, or if they have never been clicked by the same user, $\#(i, j) = 0$ and $PMI(i, j) = -\infty$. However, a negative value of PMI does not necessarily imply that the items are not related. The reason may be because the users who click i may not know about the existence of j . A common resolution to this is to replace negative values by zeros to form the PPMI matrix [3]. Elements of the PPMI matrix S are defined below:

$$S_{ij} = \max\{\widehat{PMI}(i, j), 0\} \quad (5)$$

We can see that item embedding vectors ρ and item context vectors α can be obtained by factorizing PPMI matrix S . The factorization can be performed by PMF [12].

2.3 Joint Model of Ratings and Clicks

In modeling items, we let item feature vector β_i deviate from embedding vector ρ_i . This deviation (i.e., $\beta_i - \rho_i$) accounts for the contribution of rating data in the item representation. In detail, if item i has few prior ratings, this deviation should be small; in contrast, if i has many prior ratings, this deviation should be large to allow more information from ratings to be directed toward the item representation. This deviation is introduced by letting β_i be a Gaussian distribution with mean ρ_i :

$$p(\boldsymbol{\beta}|\boldsymbol{\rho}, \sigma_\beta^2) = \prod_i \mathcal{N}(\beta_i|\rho_i, \sigma_\beta^2) \quad (6)$$

Below is the *generative process* of the model:

1. Item embedding model

(a) For each item i : draw embedding and context vectors

$$\rho_i \propto \mathcal{N}(0, \sigma_\rho^2 \mathbf{I}), \quad \alpha_i \propto \mathcal{N}(0, \sigma_\alpha^2 \mathbf{I}) \quad (7)$$

(b) For each pair (i, j) , draw S_{ij} of the PPMI matrix:

$$S_{ij} \propto \mathcal{N}(\rho_i^\top \alpha_j, \sigma_S^2) \quad (8)$$

2. Rating model

(a) For each user u : draw user feature vector and bias term

$$\theta_u \propto \mathcal{N}(0, \sigma_\theta^2 \mathbf{I}), \quad b_u \propto \mathcal{N}(0, \sigma_b^2) \quad (9)$$

(b) For each item i : draw item feature vector and bias term

$$\epsilon_i \propto \mathcal{N}(0, \sigma_\epsilon^2 \mathbf{I}), \quad c_i \propto \mathcal{N}(0, \sigma_c^2) \quad (10)$$

(c) For each pair (u, i) : draw the rating

$$R_{uj} \propto \mathcal{N}(B_{ui} + \theta_u^\top \beta_j, \sigma_R^2) \quad (11)$$

The *posterior distribution* of the model parameters given the rating matrix R , PPMI matrix S and the hyper-parameters is as follows.

$$\begin{aligned} p(\boldsymbol{\Omega}|R, S, \boldsymbol{\Theta}) &\propto P(R|\mathbf{b}, \mathbf{c}, \boldsymbol{\theta}, \boldsymbol{\beta}, \sigma_R^2) P(S|\boldsymbol{\rho}, \boldsymbol{\alpha}, \sigma_S^2) \\ &\quad \times p(\boldsymbol{\theta}|\sigma_\theta^2) p(\boldsymbol{\beta}|\boldsymbol{\rho}) p(\boldsymbol{\rho}|\sigma_\rho^2) p(\boldsymbol{\alpha}|\sigma_\alpha^2) p(\mathbf{b}|\sigma_b^2) p(\mathbf{c}|\sigma_c^2) \\ &= \prod_{(u,i) \in \mathcal{R}} \mathcal{N}(R_{ui}|\theta_u^\top \beta_i, \sigma_S^2) \prod_{(i,j) \in \mathcal{S}} \mathcal{N}(S_{ij}|\rho_i^\top \alpha_j, \sigma_S^2) \\ &\quad \times \prod_u \mathcal{N}(\theta_u|\mathbf{0}, \sigma_\theta^2) \prod_i \mathcal{N}(\beta_i|\rho_i, \sigma_\beta^2) \prod_i \mathcal{N}(\rho_i|\mathbf{0}, \sigma_\rho^2) \\ &\quad \times \prod_j \mathcal{N}(\alpha_j|\mathbf{0}, \sigma_\alpha^2) \prod_u \mathcal{N}(b_u|0, \sigma_b^2) \prod_i \mathcal{N}(c_i|0, \sigma_c^2) \end{aligned} \quad (12)$$

where $\boldsymbol{\Theta} = \{\sigma_\theta^2, \sigma_\beta^2, \sigma_\rho^2, \sigma_\alpha^2, \sigma_b^2, \sigma_c^2\}$.

2.4 Parameter Learning

Since learning the full posterior of all model parameters is intractable, we will learn the maximum a posterior (MAP) estimates. This is equivalent to minimizing the following error function.

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\Omega}) &= \frac{1}{2} \sum_{(u,i) \in \mathcal{R}} [R_{ui} - (B_{ui} + \theta_u^\top \beta_i)]^2 + \frac{\lambda}{2} \sum_{(i,j) \in \mathcal{S}} (S_{ij} - \rho_i^\top \alpha_j)^2 \\
&+ \frac{\lambda_\theta}{2} \sum_{u=1}^N \|\theta_u\|_F^2 + \frac{\lambda_\beta}{2} \sum_{i=1}^M \|\beta_i - \rho_i\|_F^2 + \frac{\lambda_\rho}{2} \sum_{i=1}^M \|\rho_i\|_F^2 \\
&+ \frac{\lambda_\alpha}{2} \sum_{j=1}^M \|\alpha_j\|_F^2 + \frac{\lambda_b}{2} \sum_{u=1}^N b_u^2 + \frac{\lambda_c}{2} \sum_{i=1}^M c_i^2
\end{aligned} \tag{13}$$

where $\lambda = \sigma_R^2/\sigma_S^2$, $\lambda_\theta = \sigma_R^2/\sigma_\theta^2$, $\lambda_\beta = \sigma_R^2/\sigma_\beta^2$, $\lambda_\rho = \sigma_R^2/\sigma_\rho^2$, and $\lambda_\alpha = \sigma_R^2/\sigma_\alpha^2$.

We optimize this function by coordinate descent; which alternatively, updates each of the variables $\{\theta_u, \beta_i, \rho_i, \alpha_j, b_u, c_i\}$ while the remaining are fixed.

For θ_u : given the current estimates of the remaining parameters, taking the partial deviation of $\mathcal{L}(\boldsymbol{\Omega})$ (Eq. 13) with respect to θ_u and setting it to zero, we obtain the update formula:

$$\theta_u = \left(\sum_{i \in \mathcal{R}_u} \beta_i \beta_i^\top + \lambda_\theta \mathbf{I}_d \right)^{-1} \sum_{i \in \mathcal{R}_u} (R_{ui} - B_{ui}) \beta_i \tag{14}$$

Similarly, we can obtain the update equations for the remaining parameters:

$$\beta_i = \left(\sum_{u \in \mathcal{R}_i} \theta_u \theta_u^\top + \lambda_\beta \mathbf{I}_d \right)^{-1} \left[\lambda_\beta \rho_i + \sum_{u \in \mathcal{R}_i} (R_{ui} - B_{ui}) \theta_u \right] \tag{15}$$

$$b_u = \frac{\sum_{i \in \mathcal{R}_u} R_{ui} - \left[\mu |\mathcal{R}_u| + \sum_{i \in \mathcal{R}_u} (c_i + \theta_u^\top \beta_i) \right]}{|\mathcal{R}_u| + \lambda_b} \tag{16}$$

$$c_i = \frac{\sum_{u \in \mathcal{R}_i} R_{ui} - \left[\mu |\mathcal{R}_i| + \sum_{u \in \mathcal{R}_i} (b_u + \theta_u^\top \beta_i) \right]}{|\mathcal{R}_i| + \lambda_c} \tag{17}$$

$$\rho_i = \left[\lambda \sum_{j \in \mathcal{S}_i} \alpha_j \alpha_j^\top + (\lambda_\beta + \lambda_\rho) \mathbf{I}_d \right]^{-1} \left(\lambda_\beta \beta_i + \lambda \sum_{j \in \mathcal{S}_i} S_{ij} \alpha_j \right) \tag{18}$$

$$\alpha_j = \left(\lambda \sum_{i \in \mathcal{S}_j} \rho_i \rho_i^\top + \lambda_\alpha \mathbf{I}_d \right)^{-1} \left(\lambda \sum_{i \in \mathcal{S}_j} S_{ij} \rho_i \right) \tag{19}$$

Computational complexity. For user vectors, as analyzed in [7], the complexity for updating N users in an iteration is $\mathcal{O}(d^2|\mathcal{R}^+| + d^3N)$. For item vector updating, we can also easily show that the running time for updating M items in an iteration is $\mathcal{O}(d^2(|\mathcal{R}^+| + |\mathcal{S}^+|) + d^3M)$. We can see that the computational complexity linearly scales with the number of users and the number of items. Furthermore, this algorithm can easily be parallelized to adapt to large scale data. For example, in updating user vectors θ , the update rule of user u is independent of other users' vectors, therefore, we can update θ_u in parallel.

2.5 Rating Prediction

We consider two cases of rating predictions: **in-matrix** prediction and **out-matrix** prediction. In-matrix prediction refers to the case where we predict the rating of user u to item i , where i has not been rated by u but has been rated by at least one of the other users; while out-matrix prediction refers to the case where we predict the rating of user u to item i , where i has not been rated by any users (i.e., only click data is available for i). The missing rating r_{ui} can be predicted using the following formula:

$$\hat{r}_{ui} \approx \mu + b_u + c_i + \theta_u^\top \beta_i \quad (20)$$

3 Empirical Study

3.1 Datasets, Competing Methods, Metric and Parameter settings

Datasets. We use three public datasets of different domains with varying sizes. The datasets are: (1) *MovieLens 1M*: a dataset of user-movie ratings, which consists of 1 million ratings in the range 1 – 5 to 4000 movies by 6000 users, (2) *MovieLens 20M*: another dataset of user-movie ratings, which consists of 20 million ratings in range 1 – 5 to 27,000 movies by 138,000 users, and (3) *Bookcrossing*: a dataset for user-book ratings, which consists of 1,149,780 ratings and clicks to 271,379 books by 278,858 users.

Since Movielens datasets contain only rating data, we artificially create the click data and rating data following [2]. Click data is obtained by binarizing the original rating data; while the rating data is obtained by randomly picking with different percentages (10%, 20%, 50%) from the original rating data. Details of datasets obtained are given in Table 2 and Table 3.

From each dataset, we randomly pick 80% of rating data for training the model, while the remaining 20% is for testing. From the training set, we randomly pick 10% as the validation set.

As discussed in Section 2, we consider two rating prediction tasks: *in-matrix* prediction and *out-matrix* prediction. In evaluating the in-matrix prediction, we ensure that all the items in the test set appear in the training set. In evaluating the out-matrix prediction, we ensure that none of the items in the test set appear in the training set.

Table 2. Datasets obtained by picking ratings from the *Movielens 1M*

Dataset	% rating picked	Density of rating matrix (%)
ML1-10	10%	0.3561
ML1-20	20%	0.6675
ML1-50	50%	1.6022

Table 3. Datasets obtained by picking ratings from *Movielens 20M*

Dataset	% rating picked	Density of rating matrix (%)
ML20-10	10%	0.0836
ML20-20	20%	0.1001
ML20-50	50%	0.2108

Competing methods. We compare EMB-MF with the following methods.

- *State-of-the-art methods* in rating predictions: PMF [12], SVD++ [8]³.
- *Item2Vec+MF*: The model was obtained by training item embedding and MF separately. First we trained an Item2Vec model [1] on click data to obtain the item embedding vectors ρ_i . We then fixed these item embedding vectors and used them as the item feature vectors β_i for rating prediction.

Metric. We used Root Mean Square Error (RMSE) to evaluate the accuracy of the models. RMSE measures the deviation between the rating predicted by the model and the true ratings (given by the test set), and is defined as follows.

$$RMSE = \sqrt{\frac{1}{|Test|} \sum_{(u,i) \in Test} (r_{ui} - \hat{r}_{ui})^2} \quad (21)$$

where $|Test|$ is the size of the test set.

Parameter settings. In all settings, we set the dimension of the latent space to $d = 20$. For PMF and SVD++, Item2Vec+MF, we used a grid search to find the optimal values of the regularization terms that produced the best performance on the validation set. For our proposed method, we explored different settings of hyper-parameters to study the effectiveness of the model.

3.2 Experimental Results

The test RMSE results for in-matrix and out-matrix prediction tasks are reported in Table 4 and Table 5.

From the experimental results, we can observe that:

- The proposed method (EMB-MF) outperforms all competing methods for all datasets on both in-matrix and out-matrix predictions.
- EMB-MF, SVD++ and Item2Vec+MF are much better than PMF, which use rating data only. This indicates that exploiting click data is a key factor to increase the prediction accuracy.

³ The results are obtain by using the LibRec library: <http://librec.net/>

Table 4. Test RMSE of in-matrix prediction. For EMB-MF, we fixed $\lambda = 1$ and used the validation set to find optimal values for the remaining hyper-parameters

Methods	ML-1m			ML-20m			Bookcrossing
	ML1-10	ML1-20	ML1-50	ML20-10	ML20-20	ML20-50	
PMF	1.1026	0.9424	0.8983	1.0071	0.8663	0.8441	2.1663
SVD++	0.9825	0.9066	0.8871	0.8947	0.8348	0.8191	1.6916
Item2Vec+MF	0.9948	0.9135	0.8984	0.9098	0.8527	0.8355	1.9014
EMB-MF (our)	0.9371	0.8719	0.8498	0.8767	0.8299	0.8024	1.6558

Table 5. Test RMSE of out-matrix prediction. For EMB-MF, we fixed $\lambda = 1$; the remaining hyper-parameters are determined using the validation set. Only *Item2Vec* is compared, because PMF and SVD++ cannot be used for out-matrix prediction

Methods	ML-1m			ML-20m			Bookcrossing
	ML1-10	ML1-20	ML1-50	ML20-10	ML20-20	ML20-50	
Item2Vec+MF	1.0986	1.0365	1.039	1.0128	0.9582	0.9784	1.7027
EMB-MF (our)	1.0312	1.0059	1.0132	0.9729	0.9422	0.9494	1.6828

- For all methods, the accuracies increase with the density of rating data. This is expected because the rating data is reliable for inferring users’ preferences.
- In all cases, the differences between EMB-MF with the competing methods are most pronounced in the most sparse subsets (ML1-10 or ML20-10). This demonstrates the effectiveness of EMB-MF on extremely sparse data.
- EMB-MF outperforms Item2Vec+MF although these two models are based on similar assumptions. This indicates the advantage of training these models jointly, rather than training them independently.

Impact of parameter λ_β . λ_β is the parameter that controls the deviation of β_i from the item embedding vector ρ_i (see Eq. 6 and Eq. 13). When λ_β is small, the value of β_i is allowed to diverge from ρ_i ; in this case, β_i mainly comes from rating data. On the other hand, when λ_β increases, β_i becomes closer to ρ_i ; in this case β_i mainly comes from click data. The test RMSE is given in Table 6.

Table 6. Test RMSE of in-matrix prediction task by the proposed method over different values of λ_β while the remaining hyper-parameters are fixed.

λ_β	0.1	1.0	10.0	20.0	50.0	100.0	1000.0
ML1-10	1.1301	0.9971	0.9318	0.9381	0.9527	0.9651	0.9924
ML1-20	1.1107	0.9963	0.8911	0.8756	0.8723	0.8769	0.8885
ML1-50	0.9798	0.9193	0.8634	0.8545	0.8512	0.8539	0.8626

We can observe that, for small values of λ_β , the model produces low prediction accuracy (high test RMSE). The reason is that when λ_β is small, the model mostly relies on the rating data which is very sparse and cannot model items

well. When λ_β increases, the model starts using click data for prediction, and the accuracy will increase. However, when λ_β reaches a certain threshold, the accuracy starts decreasing. This is because when λ_β is too large, the representations of items mainly come from click data. Therefore, the model becomes less reliable for modeling the ratings.

4 Related Work

Exploiting click data for addressing the cold-start problem has also been investigated in the literature. Co-rating [11] combines explicit (rating) and implicit (click) feedback by treating explicit feedback as a special kind of implicit feedback. The explicit feedback is normalized into the range $[0, 1]$ and is summed with the implicit feedback matrix with a fixed proportion to form a single matrix. This matrix is then factorized to obtain the latent vectors of users and items.

Wang et. al.[13] proposed Expectation-Maximization Collaborative Filtering (EMCF) which exploits both implicit and explicit feedback for recommendation. For predicting ratings for an item, which does not have any previous ratings, the ratings are inferred from the ratings of its neighbors according to click data.

The main difference between these methods with ours is that they do not have a mechanism for balancing the amounts of click data and rating data when making predictions. In our model, these amounts are controlled depending on the number of previous ratings that the target items have.

Item2Vec [1] is a neural network based model for learning item embedding vectors using *co-click* information. In [10], the authors applied a word embedding technique by factorizing the shifted PPMI matrix [9], to learn item embedding vectors from click data. However, using these vectors directly for rating prediction is not appropriate because click data does not exactly reflect preferences of users. Instead, we combine item embedding with MF in a way that allows rating data to contribute to item representations.

5 Conclusion

In this paper, we proposed a probabilistic model that exploits click data for addressing the cold-start problem in rating prediction. The model is a combination of two models: (i) an item embedding model for click data, and (ii) MF for rating prediction. The experimental results showed that our proposed method is effective in rating prediction for items with no previous ratings and also boosts the accuracy of rating prediction for extremely sparse data.

We plan to explore several ways of extending or improving this work. The first direction is to develop a full Bayesian model for inferring the full posterior distribution of model parameters, instead of point estimation which is prone to overfitting. The second direction we are planning to pursue is to develop an online learning algorithm, which updates user and item vectors when new data are collected without retraining the model from the beginning.

Acknowledgments. This work was supported by a JSPS Grant-in-Aid for Scientific Research (B) (15H02789, 15H02703).

References

1. Barkan, O., Koenigstein, N.: Item2Vec: neural item embedding for collaborative filtering. In: 26th IEEE International Workshop on Machine Learning for Signal Processing. pp. 1–6 (2016)
2. Bell, R.M., Koren, Y.: Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In: Proceedings of the 7th IEEE International Conference on Data Mining. pp. 43–52 (2007)
3. Bullinaria, J.A., Levy, J.P.: Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods*. pp. 510–526 (2007)
4. Church, K.W., Hanks, P.: Word association norms, mutual information, and lexicography. *Comput. Linguist.* pp. 22–29 (1990)
5. Gopalan, P., Hofman, J.M., Blei, D.M.: Scalable recommendation with hierarchical poisson factorization. In: Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence. pp. 326–335 (2015)
6. Gopalan, P.K., Charlin, L., Blei, D.: Content-based recommendations with poisson factorization. In: Proceedings of the 27th Advances in Neural Information Processing Systems. pp. 3176–3184 (2014)
7. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: Proceedings of the 8th IEEE International Conference on Data Mining. pp. 263–272 (2008)
8. Koren, Y.: Factorization meets the neighborhood: A multifaceted collaborative filtering model. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 426–434 (2008)
9. Levy, O., Goldberg, Y.: Neural word embedding as implicit matrix factorization. In: Proceedings of the 27th International Conference on Neural Information Processing Systems, pp. 2177–2185 (2014)
10. Liang, D., Altosaar, J., Charlin, L., Blei, D.M.: Factorization meets the item embedding: Regularizing matrix factorization with item co-occurrence. In: Proceedings of the 10th ACM Conference on Recommender Systems. pp. 59–66 (2016)
11. Liu, N.N., Xiang, E.W., Zhao, M., Yang, Q.: Unifying explicit and implicit feedback for collaborative filtering. In: Proceedings of the 19th ACM International Conference on Information and Knowledge Management. pp. 1445–1448 (2010)
12. Mnih, A., Salakhutdinov, R.R.: Probabilistic matrix factorization. In: 20th Advances in Neural Information Processing Systems. pp. 1257–1264 (2008)
13. Wang, B., Rahimi, M., Zhou, D., Wang, X.: Expectation-maximization collaborative filtering with explicit and implicit feedback. In: The 16th Pacific-Asia Conference on Knowledge Discovery and Data Mining. pp. 604–616 (2012)
14. Wang, C., Blei, D.M.: Collaborative topic modeling for recommending scientific articles. In: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 448–456 (2011)
15. van den Oord, A., Dieleman, S., Schrauwen, B.: Deep content-based music recommendation. In: Proceedings of the Advances in Neural Information Processing Systems 26. pp. 2643–2651 (2013)