SIFN: A Sentiment-aware Interactive Fusion Network for Review-based Item Recommendation

Kai Zhang[1], Hao Qian[3], Qi Liu[1,2,*], Zhiqiang Zhang[3], Jun Zhou[3], Jianhui Ma[1], Enhong Chen[1,2]

[1]Anhui Province Key Lab of Big Data Analysis and Application, School of Data Science,
University of Science and Technology of China; [3]Ant Group, Hangzhou, China
[2]School of Computer Science and Technology, University of Science and Technology of China
kkzhang0808@mail.ustc.edu.cn, {qiliuql, jianhui, cheneh}@ustc.edu.cn
{qianhao.qh, lingyao.zzq, jun.zhoujun}@antgroup.com

ABSTRACT

Recent studies in recommender systems have managed to achieve significantly improved performance by leveraging reviews for rating prediction. However, despite being extensively studied, these methods still suffer from some limitations. First, previous studies either encode the document or extract latent sentiment via neural networks, which are difficult to interpret the sentiment of reviewers intuitively. Second, they neglect the personalized interaction of reviews with user/item, i.e., each review has different contributions when modeling the sentiment preference of user/item.

To remedy these issues, we propose a **S**entiment-aware **I**nteractive **F**usion **N**etwork (**SIFN**) for review-based item recommendation. Specifically, we first encode user/item reviews via BERT and propose a light-weighted sentiment learner to extract semantic features of each review. Then, we propose a sentiment prediction task that guides the sentiment learner to extract sentiment-aware features via explicit sentiment labels. Finally, we design a rating prediction task that contains a rating learner with an interactive and fusion module to fuse the identity (i.e., user and item ID) and each review representation so that various interactive features can synergistically influence the final rating score. Experimental results on five real-world datasets demonstrate that the proposed model is superior to state-of-the-art models.

KEYWORDS

Review-based Recommendation; Interactive Fusion Network

* Corresponding Author.

1 INTRODUCTION

Recommender Systems (RS) have been widely adopted by many platforms, e.g. Amazon and Yelp, thanks to their capability of filtering information based on the user interests [5, 16]. In these platforms, user reviews of items contain rich semantic information. Therefore, to better estimate user's rating, review-based recommendation has gained wide attention in both academia [1] and industry [7, 12].

In the literature, there are many efforts for this recommendation problem, especially on the rating prediction task. Among them, the most representative methods are based on Matrix Factorization (MF) model, which decomposes the user-item rating matrix into two matrices corresponding to user and item features. However, these methods represent user and item information only based on ratings, which suffer from sparsity issue. To mitigate this problem, many RS have been proposed to exploit the semantic information from user reviews [1, 2, 13, 14, 18]. Among these review-based methods, the earlier works adopt MF to extract the semantics from user/item reviews, such as PMF[8] and ConvMF+ [6]. Recently, the focus of research shifts to learn latent features from reviews via neural network methods. Among them, DeepCoNN [18] applies convolution to process the reviews to learn user and item representations. D-Attn [12] leverages global and local attention to select decisive words in the review documents for rating prediction. NARRE [2] is a review retrieval model that adopts an attention mechanism to select appropriate reviews for items. Moreover, CARP [7] utilizes a sentiment capsule network to estimate user-item ratings and provide interpretability with a fine-grained manner.

Although these works achieved significant performance improvement in the review-based recommendation, they still suffer from intrinsic issues. First, most previous works directly encode reviews to extract implicit semantics while largely ignoring the explicit sentiment polarity of reviews [12, 14], which carry the user attitudes and preferences (i.e., which kinds of item user may like or dislike). Therefore, implicitly mining the semantic information of reviews may lead to sub-optimal prediction because the reviews' sentiment label has not been applied to the training process [4, 10, 14, 17]. Second, in the rating prediction scenario, reviews' sentiment polarity plays distinct roles in representing the overall sentiment preference of the user/item. For instance, the impact of a negative review may be far more significant than a positive review for the user/item's semantic preference representation. Therefore, it is necessary to weigh and fuse the interactions between each review and user/item features for the final rating prediction.
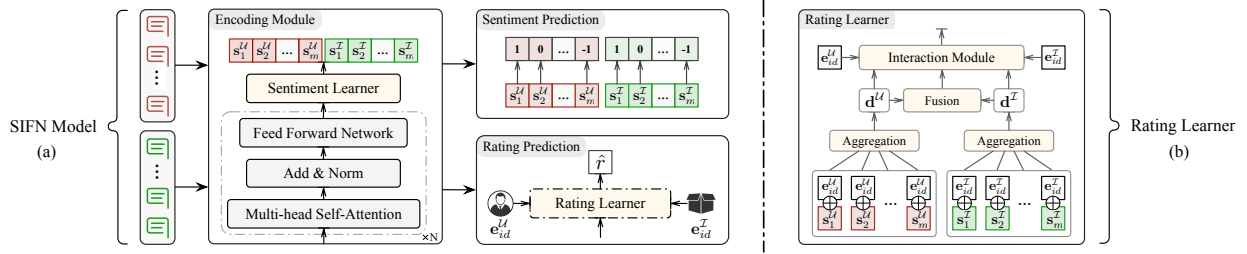
Figure 1: (a) The overall architecture of SIFN model; (b) the rating learner which can exploit interactive fusion knowledge.

To address these aforementioned issues, we propose a **S**entiment-aware **I**nteractive **F**usion **N**etwork (**SIFN**), which includes a sentiment prediction task to exploit the sentiment polarity from each review, and a rating prediction task to estimate user rating of items. Specifically, we first utilize the pre-trained BERT to encode text reviews and introduce a light-weighted sentiment learner to mine the sentiment-aware features. Then, we design a sentiment prediction task that explicitly extracts important sentiment features in each review for user/item. Finally, in the rating prediction task, we develop a rating leaner which contains two novel operations (i.e., interactive & fusion) to fuse the identity representation of user/item and review representation so that various interactive features can synergistically affect the final rating prediction. The main contributions of this paper are as follows: 1) We highlight the explicit sentiment polarity in each review, and focus on modeling the multiple feature interactions between each review and user/item. 2) We propose a novel **S**entiment-aware **I**nteractive **F**usion **N**etwork (**SIFN**) model with two main components, Sentiment Leaner and Rating Learner. 3) We conduct extensive experiments on five datasets so that the results demonstrate the effectiveness of the proposed method.

2 THE PROPOSED MODEL

Figure 1 illustrates the overall architecture of SIFN, which mainly consists of three components. The details of each component will be explained in the following sections.

2.1 The Encoding Module

In this section, we present how to extract deep semantic representations of reviews progressively. Since the semantic learning process of reviews for user and item are quite similar, we mainly introduce user reviews' semantic representation for space saving.

BERT Encoding. Given a user review $x^{\mathcal{U}} = (w_1^{\mathcal{U}}, w_2^{\mathcal{U}}, ..., w_l^{\mathcal{U}}) \in \mathbb{R}^l$ which is formed by merging the total l words in a review written by a user. For word embedding, we adopt the pre-trained BERT [3] to encode each word and use the last hidden state of the pre-trained BERT for word representations [15]. The formulation is as follows:

$$\mathrm{R}^{\mathcal{U}} = (\mathrm{e}_1^{\mathcal{U}}, \mathrm{e}_2^{\mathcal{U}}, ..., \mathrm{e}_l^{\mathcal{U}}) = \mathrm{BERT}(w_1^{\mathcal{U}}, w_2^{\mathcal{U}}, ..., w_l^{\mathcal{U}}), \quad (1)$$

where $\mathrm{R}^{\mathcal{U}} \in \mathbb{R}^{l \times k}$ is the embedding of a user review, $\mathrm{e}_l^{\mathcal{U}} \in \mathbb{R}^k$ is the embedding of the l-th word, where k is the word embedding size. Similarly, an item review can be mapped as $\mathrm{R}^{\mathcal{I}} = (e_1^{\mathcal{I}}, e_2^{\mathcal{I}}, ..., e_l^{\mathcal{I}}) \in \mathbb{R}^{l \times k}$. With the powerful pre-trained models, the general semantic knowledge of the text reviews can be fully extracted.

Sentiment Learner. As each word in a review carries distinct semantic and sentimental information, which is crucial to estimate user rating and sentiment polarity. Therefore, to better suit the downstream tasks, we devise a sentiment learner to adaptively underline the informative words. Specifically, we employ the attention mechanism, in which the attention score of i-th word is calculated as:

$$\alpha_i = \frac{\exp(\tanh(\mathrm{W}_a \mathrm{e}_i^{\mathcal{U}} + \mathrm{b}_a))}{\sum_{i'=1}^{l} \exp(\tanh(\mathrm{W}_a \mathrm{e}_{i'}^{\mathcal{U}} + \mathrm{b}_a))}. \quad (2)$$

W_a and b_a are weight matrix and bias, respectively. α_i is attention score of i-th word in a review. Therefore, words with smaller attention score are less important. Further, with a weighted sum pooling over the words in a review, the effects of uninformative words are diminished and the semantic representation of the user review $\mathrm{s}^{\mathcal{U}} \in \mathbb{R}^k$ is calculated as:

$$\mathrm{s}^{\mathcal{U}} = \sum_{i=1}^{l} \alpha_i \mathrm{e}_i^{\mathcal{U}}, \quad (3)$$

Through the encoding module, we obtain each review's representation of the user, i.e., $\mathrm{s}^{\mathcal{U}}$, via aggregating feature vectors of the words. Similarly, we use the same method to generate each review's representation, i.e., $\mathrm{s}^{\mathcal{I}} \in \mathbb{R}^k$, for the item.

2.2 Sentiment Prediction

Since the sentiment polarity of review carries the intrinsic user preference toward item, we design a sentiment prediction task for each review to learn sentiment features explicitly (e.g., $\mathrm{s}^{\mathcal{U}}$ and $\mathrm{s}^{\mathcal{I}}$). Note that, the ground-truth is obtained by converting the user-item rating to sentiment label with a threshold of 3, i.e., we category them into positive (i.e., higher than 3 stars), negative (i.e., lower than 3 stars) or neural (i.e., equal to 3 stars).

Hereby, we utilize a cross entropy loss function to estimate the sentiment polarity of each user and item review as follows:

$$L_s = -\frac{1}{|O|} \sum_{(\mathcal{U},\mathcal{I}) \in O} \{\frac{1}{m} \sum_{j=1}^{m} (\sum_{t=1}^{C} y_j^t \log(\hat{y}_j^t))\}, \quad (4)$$

where $\hat{y}_j^t = softmax(\mathrm{s}_j^{\mathcal{U}})$ is the predicted sentiment polarity of the j-th user review. Similarly, we can get sentiment prediction of item review $\hat{y}_j^t = softmax(\mathrm{s}_j^{\mathcal{I}})$. O denotes the set of observed user-item rating pairs. C is the number of sentiment polarity categories and m is the number of reviews for a user/item.

2.3 Rating Prediction

In this subsection, we apply our SIFN model for the rating prediction task. To this end, we propose a rating learner to comprehensively extract the latent interactions among the identity representation of user/item and review representation, as shown in figure 1 (b).

To begin with, there is user ID feature, which is encoded into low-dimensional vector denoted as $e_{id}^{\mathcal{U}} \in \mathbb{R}^k$ through:

$$e_{id}^{\mathcal{U}} = W_{id}^{\mathcal{U}} \cdot x_{id}^{\mathcal{U}}, \tag{5}$$

where $W_{id}^{\mathcal{U}} \in \mathbb{R}^{n^{\mathcal{U}} \times k}$ is weight matrix and $n^{\mathcal{U}}$ is the total number of user ID features.

Rating Learner. As previous work states [2], different reviews do not contribute equally to the preference of user/item. Besides, user also carries distinct style when writing reviews. For example, an optimistic user is more likely to write positive reviews while a cranky user tends to write something even rude. Therefore, in order to better model the user/item's personalized preference in different reviews, it is indispensable to incorporate both text review and user/item features. Hereby, we concatenate them as follows:

$$o_j^{\mathcal{U}} = W_o \cdot \mathrm{concat}(s_j^{\mathcal{U}}; e_{id}^{\mathcal{U}}), \tag{6}$$

where $W_o \in \mathbb{R}^{k \times 2k}$ is a projection weight matrix.

1) Aggregation. Then, we utilize the user-aware review representation $o_j^{\mathcal{U}} \in \mathbb{R}^k$ to adaptively select important reviews with attention function which is formulated as:

$$d^{\mathcal{U}} = \sum_{j=1}^{m} \frac{\exp\left(\tanh\left(W_o o_j^{\mathcal{U}} + b_o\right)\right)}{\sum_{j'=1}^{m} \exp\left(\tanh\left(W_o o_{j'}^{\mathcal{U}} + b_o\right)\right)} o_j^{\mathcal{U}}, \tag{7}$$

where $d^{\mathcal{U}} \in \mathbb{R}^k$ is the user review representation with aggregated information from all user reviews under the influence of user features. Similarly, we can get the embedding vector of item ID denoted as $e_{id}^{\mathcal{I}} \in \mathbb{R}^k$ and the item review representation $d^{\mathcal{I}} \in \mathbb{R}^k$.

2) Fusion Network. So far, the user/item reviews representation (i.e., $o_j^{\mathcal{U}}$ and $o_j^{\mathcal{I}}$), that integrates the features of identity and reviews, are learned from separate hierarchical representation modules without interactions. Since we aim at learning the user-item rating, it is necessary to extract the mutual influence among them. However, as user and item features may carry different characteristics in different spaces, we design a novel fusion network that projects features into a shared hidden space before learning feature interactions as:

$$f = d^{\mathcal{U}} W_f d^{\mathcal{I}}, \tag{8}$$

where $f \in \mathbb{R}^k$ is the fusion representation, $W_f \in \mathbb{R}^{k \times k}$ is the weight matrix that maps user and item features to the same latent space.

3) Interactive Network. In addition, to underline the contribution of the different extracted representations (e.g., $d^{\mathcal{U}}$, $d^{\mathcal{I}}$ and f), we further design an interactive network so that the user-item rating estimation is fully explored as follows:

$$p = \left(d^{\mathcal{U}} + e_{id}^{\mathcal{U}}\right) \odot \left(d^{\mathcal{I}} + e_{id}^{\mathcal{I}}\right) + Wf + b, \tag{9}$$

where W and b are weight matrix and bias. $\odot$ denotes the element-wise product of vectors. $p \in \mathbb{R}^k$ is the user-item preference representation with aggregated interactions from reviews and identity.

Table 1: Statistics of the Amazon datasets.

Dataset	# Users	# Items	# Ratings	Density
Music Instruments	1,429	900	10,261	0.798%
Office Products	4,905	2,420	53,228	0.448%
Digital Music	5,540	3,568	64,664	0.327%
Tools	16,638	10,217	134,345	0.079%
Video Games	24,303	10,672	213,577	0.089%

Through the above learning processes, the predicted rating of user toward item $\hat{r}$ is obtained by a linear projection and the regression loss function for rating prediction is as:

$$\hat{r} = w_r p + b^{\mathcal{U}} + b^{\mathcal{I}}; \quad L_r = \frac{1}{|O|} \sum_O (r - \hat{r})^2, \tag{10}$$

where w_r is the weight parameter. $b^{\mathcal{U}}$ and $b^{\mathcal{I}}$ are biases for user and item, respectively. r is the ground truth user-item rating.

Finally, we utilize the joint learning process to optimize both objectives with a hyper-parameters λ as:

$$L = L_r + \lambda \cdot L_s. \tag{11}$$

3 EXPERIMENT

3.1 Dataset

We choose five Amazon datasets (i.e., Music Instruments, Office Products, Digital Music, Tools, Video Games) from [7] to conduct our experiments. We preprocessed it to ensure that all users and items have at least one rating five reviews in our experiment. The dataset consists of numerous data samples on which we follow a randomized 80:10:10 train/test/validation split. For concrete information, please refer to Table 1.

3.2 Baseline methods

To show the performance of our proposed model, we compare SIFN with the state-of-the-art (SOTA) models. The benchmarks are:

- **MF-based**: PMF [8] models the latent factors by introducing Gaussian distribution. ConvMF+ [6] incorporates convolutional neural network into Matrix Factorization.
- **Neural-based**: DeepCoNN [18], D-Attn [12], NARRE [2], CARP [7]. These methods have been introduced in section 1.

In our SIFN model, we use BERT encoding from Hugging Face [1]. User and item embedding size is set to 16. The batch size is 100, learning rate is 0.001, dropout rate is 0.2, and λ is tuned amongst [0.1,1,10]. For baseline methods, we follow the hyper-parameter configurations in their papers. Following previous works [7, 12], we utilize Mean Squared Error (MSE) as the evaluation metric and select Adam as optimizer for all models.

3.3 Experimental Results

In this section, we evaluate the performance of SIFN model on Amazon datasets along with detailed comparison of results in Table 2. The major results are summarized as follows: (1) MF-based methods (e.g., PMF and ConvMF+) consistently fall behind other methods, which indicates the limitations of Matrix Factorization to learn

---

[1] https://huggingface.co/transformers/

Table 2: Experimental results in Amazon datasets. Percentage in ( ) denotes the relative improvement of SIFN over the baseline method, which is achieved with paired t-tests at the significance level of 0.01. We underline the best performed baseline.

| Methods	Music Instruments	Office Products	Digital Music	Tools	Video Games	Average
PMF [8]	1.398(+45.7%)	1.092(+35.7%)	1.206(+33.7%)	1.566(+39.3%)	1.672(+37.4%)	1.386(+38.6%)
ConvMF+ [6]	0.991(+23.4%)	0.960(+26.9%)	1.084(+26.3%)	1.240(+23.4%)	1.449+27.7%)	1.145(+25.7%)
DeepCoNN [18]	0.814(+6.76%)	0.860(+18.4%)	1.058(+24.5%)	1.061(+10.5%)	1.145(+8.56%)	0.988(+13.9%)
D-Attn [12]	0.982(+22.7%)	0.825(+14.9%)	0.911(+12.3%)	1.043(+8.92%)	1.144(+8.48%)	0.981(+13.3%)
NARRE [2]	0.803(+5.48%)	0.848(+17.2%)	0.898(+11.0%)	1.029(+7.68%)	1.129(+7.26%)	0.941(+9.6%)
CARP [7]	0.773(+1.81%)	0.719(+2.36%)	0.820(+2.56%)	0.960(+1.04%)	1.084(+3.41%)	0.872(+2.4%)
SIFN	**0.759**	**0.702**	**0.799**	**0.950**	**1.047**	**0.851**
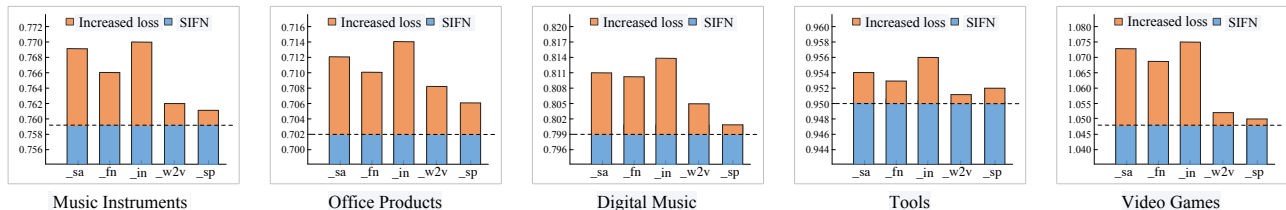
Figure 2: Results from variants of SIFN. In each plot, the x-axis denotes the name of variants, in which the prefix SIFN is omitted for space saving. The y-axis denotes the value of evaluation loss. The blue bar represents the loss from SIFN while the stacked orange part denotes increased loss, which indicates a worse performance of variant than SIFN.

semantic information from sparse rating dataset. (2) Neural-based methods (e.g., DeepCoNN, D-Attn, NARRE, CARP) outperform MF-based ones by a large margin, which validates the powerful feature extraction capability of neural networks. Among them, CARP performs best, which suggests it's helpful to encode viewpoint of user and sentiment aspect of item for rating prediction. (3) Furthermore, our proposed SIFN model still outperforms CARP by 1.81%~3.41%, which demonstrates the superiority of the well designed interactive& Fusion module that extracts the mutual influence among user and item features and the effectiveness of the sentiment prediction task that attends the sentimental phrases in text reviews.

3.4 Ablation Studies

To explore the impact of each component of SIFN, multiple ablation studies are carried out by removing one sub-module at a time.

- **SIFN_sa**: replaces the sentence attention with a simple average sum pooling over all the reviews.
- **SIFN_fn**: removes the fusion network so that user and item features are disentangled without explicit interactions.
- **SIFN_in**: replaces the interactive network with commonly used Factorization Machine (FM) [11] to estimate the ratings.
- **SIFN_w2v**: replaces the BERT encoding of text reviews with commonly used pre-trained word embedding GloVe [9].
- **SIFN_sp**: removes the sentiment prediction task so that the model focuses on user-item rating prediction.

We report the results of ablation studies in Figure 2. Specifically, the performance of **SIFN_sa** with average pooling drops as it just assumes every review contributes equally for user-item rating, which overlooks the impact of various user/item information. Without the fusion network, the performance of **SIFN_fn** also declines since
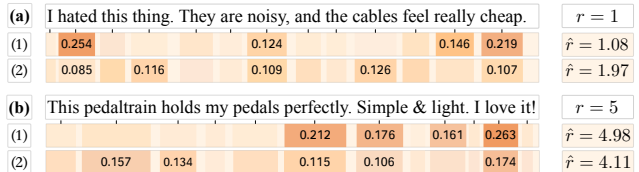

Figure 3: Attention visualization of the different words. The word attention value from SIFN and SIFN_sp are associated with the row (1), row (2) respectively in both examples.

the interactions of user and item reviews are not fully exploited. As user and item features carry different characteristics, it is inefficient to perform second-order operations in the same space with FM, which leads to the decrease of the performance of **SIFN_in**. Not surprisingly, without BERT encoding, the word embedding in **SIFN_w2v** is incapable of representing deep semantic information of text reviews. Without sentiment prediction in **SIFN_sp**, there is no supervision towards attending sentiment-aware words in reviews, which are vital for the rating prediction.

3.5 Case Study

As shown in Figure 3, we randomly sample two reviews of a user from the Music Instrument dataset and visualize the attention results to explain the capability of SIFN in extracting sentiment knowledge. Specifically, in review (a), SIFN aligns the sentiment words, e.g., "hated" and "cheap", with a rating of 1.08, which is consistent with actual value. In contrast, we utilize the attention results from **SIFN_sp** as described in Section 3.4, which overlooks these sentiment words thus predicting a relatively higher rating. Similarly, in review (b), SIFN accurately predicts a rating of 4.98 by extracting

the sentiment words, e.g., "love" and "perfectly", while **SIFN_sp** is incapable of achieving this. These examples demonstrate that SIFN is effective in extracting the latent semantics of reviews and interpret the corresponding sentiment, which is helpful for the final rating prediction.

4 CONCLUSIONS

In this paper, we proposed a novel Sentiment-aware Interactive Fusion Network (SIFN) model for review-based item recommendation. Specifically, we first employed the encoding module which contains BERT encoding and a sentiment learner to learn sentiment-aware features of each review sentence. Then, we designed a sentiment prediction task to guide the learn process of sentiment features. Finally, we developed a rating learner for the final rating prediction, which contains two novel operations (i.e., interactive & fusion) to fuse the identity features of user/item and review representation. Extensive experiments on five public datasets demonstrate the effectiveness of our proposed method.

5 ACKNOWLEDGMENT

REFERENCES

[1] Yang Bao, Hui Fang, and Jie Zhang. 2014. Topicmf: Simultaneously exploiting ratings and reviews for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 28.

[2] Chong Chen, Min Zhang, Yiqun Liu, and Shaoping Ma. 2018. Neural attentional rating regression with review-level explanations. In *Proceedings of the 2018 World Wide Web Conference*. 1583–1592.

[3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[4] Xiangnan He, Tao Chen, Min-Yen Kan, and Xiao Chen. 2015. Trirank: Review-aware explainable recommendation by modeling aspects. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. 1661–1670.

[5] Min Hou, Le Wu, Enhong Chen, Zhi Li, Vincent W Zheng, and Qi Liu. 2019. Explainable fashion recommendation: A semantic attribute region guided approach. *arXiv preprint arXiv:1905.12862* (2019).

[6] Donghyun Kim, Chanyoung Park, Jinoh Oh, Sungyoung Lee, and Hwanjo Yu. 2016. Convolutional matrix factorization for document context-aware recommendation. In *Proceedings of the 10th ACM conference on recommender systems*. 233–240.

[7] Chenliang Li, Cong Quan, Li Peng, Yunwei Qi, Yuming Deng, and Libing Wu. 2019. A capsule network for recommendation and explaining what you like and dislike. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 275–284.

[8] Andriy Mnih and Russ R Salakhutdinov. 2007. Probabilistic matrix factorization. *Advances in neural information processing systems* 20 (2007), 1257–1264.

[9] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.

[10] Zhaochun Ren, Shangsong Liang, Piji Li, Shuaiqiang Wang, and Maarten de Rijke. 2017. Social collaborative viewpoint regression with explainable recommendations. In *Proceedings of the tenth ACM international conference on web search and data mining*. 485–494.

[11] Steffen Rendle. 2012. Factorization Machines with LibFM. 3, 3 (2012).

[12] Sungyong Seo, Jing Huang, Hao Yang, and Yan Liu. 2017. Interpretable convolutional neural networks with dual local and global attention for review rating prediction. In *Proceedings of the eleventh ACM conference on recommender systems*. 297–305.

[13] Yunzhi Tan, Min Zhang, Yiqun Liu, and Shaoping Ma. 2016. Rating-boosted latent topics: Understanding users and items with ratings and reviews.. In *IJCAI*, Vol. 16. 2640–2646.

[14] Yi Tay, Anh Tuan Luu, and Siu Cheung Hui. 2018. Multi-pointer co-attention networks for recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2309–2318.

[15] Kai Zhang, Qi Liu, Hao Qian, Biao Xiang, Qing Cui, Jun Zhou, and Enhong Chen. 2021. EATN: An Efficient Adaptive Transfer Network for Aspect-level Sentiment Analysis. *IEEE Transactions on Knowledge and Data Engineering* (2021).

[16] Kai Zhang, Hao Qian, Qing Cui, Qi Liu, Longfei Li, Jun Zhou, Jianhui Ma, and Enhong Chen. 2021. Multi-Interactive Attention Network for Fine-grained Feature Learning in CTR Prediction. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 984–992.

[17] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 83–92.

[18] Lei Zheng, Vahid Noroozi, and Philip S Yu. 2017. Joint deep modeling of users and items using reviews for recommendation. In *Proceedings of the tenth ACM international conference on web search and data mining*. 425–434.