Deep heterogeneous autoencoders for Collaborative Filtering

Li, Tianyu; Ma, Yukun; Xu, Jiu; Stenger, Björn; Liu, Chen; Hirate, Yu

2018

Li, T., Ma, Y., Xu, J., Stenger, B., Liu, C., & Hirate, Y. (2018). Deep heterogeneous autoencoders for Collaborative Filtering. Proceedings of the 2018 IEEE International Conference on Data Mining (ICDM), 1164-1169. doi:10.1109/icdm.2018.00153

<https://hdl.handle.net/10356/144026>

<https://doi.org/10.1109/ICDM.2018.00153>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/ICDM.2018.00153>

Downloaded on 26 Apr 2023 18:18:54 SGT

Deep Heterogeneous Autoencoders for Collaborative Filtering

Tianyu Li*, Yukun Ma[†], Jiu Xu*, Björn Stenger*, Chen Liu*, Yu Hirate*

*Rakuten Institute of Technology

[†]Nanyang Technological University

Abstract—This paper leverages heterogeneous auxiliary information to address the data sparsity problem of recommender systems. We propose a model that learns a shared feature space from heterogeneous data, such as item descriptions, product tags and online purchase history, to obtain better predictions. Our model consists of autoencoders, not only for numerical and categorical data, but also for sequential data, which enables capturing user tastes, item characteristics and the recent dynamics of user preference. We learn the autoencoder architecture for each data source independently in order to better model their statistical properties. Our evaluation on two *MovieLens* datasets and an e-commerce dataset shows that mean average precision and recall improve over state-of-the-art methods.

Index Terms—Deep Autoencoder, Heterogeneous Data, Shared Representation, Sequential Data Modeling, Collaborative Filtering

I. INTRODUCTION

Although Collaborative Filtering (CF) techniques achieve good performance in many recommender systems [1], their performance degrades significantly when historical data is sparse. In order to alleviate this problem, features from auxiliary data sources that reflect user preference have been extracted [2], [5], as shown in Fig. 1. How to represent data from different sources is still a research problem, and it has been shown that the representation itself substantially impacts performance [6], [20]. Recently, representation learning that automatically discovers hidden factors from raw data has become a popular approach to remedy the data sparsity issue of recommender systems [10], [14].

Many online shopping platforms gather not only user profiles and item descriptions, but various other types of data, such as product reviews, tags and images. Recent research has added textual and visual information to recommender systems [3], [4]. However, in many cases sequential data, such as user purchase and browsing history, which carries information about trends in user tastes, have largely been neglected in CF-based recommender systems.

In this paper we propose Deep Heterogeneous Autoencoders (DHA) for Collaborative Filtering to combine information from multiple domains. We use Stacked Denoising Autoencoders (SDAE) to extract latent features from non-sequential data, and Recurrent Neural Network Encoder-Decoders (RNNE) to extract features from sequential data. The model is able to capture both user preferences and potential shifts of interest over time. Each data source is

Fig. 1: **Auxiliary information usage in recommender systems.** Item descriptions and user profiles are typically being used for feature extraction to alleviate the data sparsity problem. Our proposal also leverages sequential data, such as purchase histories, to reflect user preferences.

modeled using an independent encoder-decoder mechanism. Different encoders can have different number of hidden layers and an arbitrary number of hidden units in order to deal with the intrinsic difference of data sources. For instance, user demographic data and item content are typically categorical, while user comments or item tags are textual. After pre-processing, such as one hot encoding, bag-of-words and word2vec computation, representation vectors are on a different level of abstraction. Owing to its flexible structure, our model is able to learn suitable latent feature vectors for each component. These local representations from each data source are joined to form a shared feature space, which couples the joint learning of the representation from heterogeneous data and the collaborative filtering of user-item relationships.

The contributions of this paper are summarized as follows:

- 1) A method for modeling both static and sequential data in a consistent way for recommender systems in order to capture the trend in user tastes, and
- 2) Adaptation of the autoencoder architecture to accurately model each data source by considering their distinct abstraction levels.

We show improvements in terms of mean average precision and recall on three different datasets.

II. RELATED WORK

A. Incorporating side information into recommender systems

In order to improve recommendation performance, research has been focusing on using side information, such as user pro-

Fig. 2: Deep Heterogeneous Autoencoders and the integration with collaborative filtering. The proposed model extracts a shared feature space from multiple sources of auxiliary information. It models non-sequential and sequential data to capture user preferences, item properties as well as temporal dynamics. It adopts independent encoder-decoder architectures for different data sources in order to better model their statistical properties. The product of $U \in \mathbb{R}^{m \times d}$ and $V \in \mathbb{R}^{n \times d}$ approximates the user-item interaction matrix.

files and reviews [3], [5]. In particular, deep learning models have been widely studied [13], [15]. AutoRec first proposed the use of autoencoders for recommender systems [17]. In more recent work, representations are learned via stacked autoencoders (SAE), and fed into conventional CF models, either loosely or tightly coupled [7], [18]. Deep models that integrate autoencoders into collaborative filtering have shown state-of-the-art performance.

B. Recurrent Neural Network Encoder-Decoder

Recurrent neural networks (RNNs) process sequential data one element at each step to capture temporal dynamics. The encoder-decoder mechanism was initially applied to RNN for machine translation [11]. Recently, RNN encoder-decoders (RNNE) have been used to learn features from a series of actions and have successfully been applied in other areas. It was shown that Long Short-Term Memory (LSTM) networks have the ability to learn on data with long range temporal dependencies, and we adopt LSTMs for modeling sequential data.

III. DEEP HETEROGENEOUS AUTOENCODERS FOR COLLABORATIVE FILTERING

A. Overview

We propose a model that learns a joint representation from heterogeneous auxiliary information to mitigate the data sparsity problem of recommender systems. SDAEs are applied to numerical and categorical data for modeling the static tastes

of users for items. We use RNNEs to extract features from sequential data to reveal interest shifts over time.

The model adopts an independent autoencoder architecture for each data source since the inputs are generally on a different level of abstraction, see Fig. 2 for an overview. In order to discover the distinct statistical properties of every data source, our model takes the existing disparity of input abstraction levels into consideration, and applies autoencoders to each source independently by allowing distinct hidden layer numbers and arbitrary hidden units at every layer.

B. Deep Heterogeneous Autoencoders

We define each source of auxiliary data as a component indexed by $c \in \{1, \dots, C\}$. S_c denotes the input of component c . We pre-process non-sequential data like textual item descriptions by generating fixed-length embedding vectors. For sequential data, an embedding vector is learned for every time step after tokenization. We separately describe the encoding-decoding outputs of the above two types of embedding vectors.

As shown in Fig. 2, SDAE is applied to fixed-length embedding vectors. Each component encoder takes the input S_c , generates a corrupted version of it, $\hat{S}_c$, and the first layer maps it to a hidden representation h_c , which captures the main factors of variation in the input data distribution [8], [9]. More importantly, the number of component hidden layers in our model can differ from each other. The architecture is unique for each data source, where the number of layers of component c is denoted as L_c . The representation at every layer is $S_{c,l}$. For the encoder of each component, given $l_c \in \{1, \dots, L_c/2\}$ and $c \in C$, the hidden representation $h_{c,l}$ is derived as:

$$h_{c,l} = f(W_{c,l}h_{c,l-1} + b_{c,l}). \quad (1)$$

The decoder reconstructs the data at layer L as follows:

$$\hat{S}_c = g(W'_{c,L}h_{c,L} + b'_{c,L}). \quad (2)$$

The proposed model leverages sequential data by using two LSTMs for encoding and decoding one sequential data source. Specifically, the encoder reads a sequence with T time steps. At the last time step, the hidden state h_T is mapped to a context vector c , as a summary of the whole input sequence [11]. The decoder generates the output sequence by predicting the next action y_t given h_t . Both y_t and h_t are also conditioned on y_{t-1} and the context vector c .

To combine them, as shown in Fig. 2, the first part of our model encodes all components to generate hidden representations $S_{c,L_c/2}$ of non-sequential data and h_T of sequential data across all sources. These are merged to generate a joint latent representation, denoted as $h_{+,0}$. Analogous to the hidden layers of each component, the fusion model can have multiple hidden layers, the total number denoted as L_+ . The representation of the first fusion hidden layer is

$$h_{+,0} = f\left(\sum_{c \in C} W_{c,+}h_{c,L_c/2} + b_{+,0}\right). \quad (3)$$

The first hidden layer $h_{+,0}$ of the fused model is fed into the collaborative filtering model. After joint training, $h_{+,0}$ is the latent vector to generate recommendation results.

C. DHA-based Collaborative Filtering

All data is fed into two DHAs for users and items, respectively. Fig. 2 shows the process for items, and it is analogous for user data. Let $R \in \mathbb{R}^{m \times n}$ denote the rating matrix of users to items, $S_c^{(u)}$ being the component c input for users and $S_c^{(v)}$ that for items. Then, $h_{+,0}^{(u)}$ and $h_{+,0}^{(v)}$ are the latent factors. The loss function of the proposed DHA based collaborative filtering is defined as:

$$\begin{aligned} L = & \sum_{i,j} c_{i,j} (r_{i,j} - u_i v_j)^2 + \lambda_f (\sum_i \|u_i\|^2 + \sum_j \|v_j\|^2) \\ & + \lambda_m \sum_{c \in C_u} \text{loss}(S_c^{(u)}, \bar{S}_c^{(u)}) + \lambda_n \sum_{c \in C_i} \text{loss}(S_c^{(v)}, \bar{S}_c^{(v)}) \\ & + \lambda_u \sum_i \|u_i - h_{+,0}^{(u)}\|^2 + \lambda_v \sum_j \|v_j - h_{+,0}^{(v)}\|^2 \\ & + \lambda_w (\sum_c \sum_l (\|W_{c,l}^{(u)}\|^2 + \|b_{c,l}^{(u)}\|^2) \\ & + \sum_c \sum_l (\|W_{c,l}^{(v)}\|^2 + \|b_{c,l}^{(v)}\|^2)). \end{aligned} \quad (4)$$

The loss function includes reconstruction costs of user and item information sets, the error to predict $r_{i,j}$, and the approximation error between latent factor vectors of feature learning and collaborative filtering. The loss function is minimized to obtain parameters for the DHAs and the CF model. The mean squared error and the negative log-likelihood are used as cost functions for non-sequential and sequential data, separately. We use λ_m , λ_n , λ_u and λ_v to balance losses between users and items, λ_f , and λ_w to regularize the weight matrix and bias vectors.

D. Parameter learning

We apply coordinate descent to alternate the optimization between representation learning of heterogeneous data and user-item interaction, similar to [7], [16]. Given W s and b s, the gradients of the loss function L with respect to u_i and v_j are computed and set to 0, leading to the following updates:

$$u_i \leftarrow (V^T C_i V + \lambda_f I + \lambda_u I)^{-1} (V^T C_i R_i + \lambda_u h_{+,0}^{(u)}), \quad (5)$$

$$v_j \leftarrow (U^T C_j U + \lambda_f I + \lambda_v I)^{-1} (U^T C_j R_j + \lambda_v h_{+,0}^{(v)}), \quad (6)$$

where $U \in \mathbb{R}^{m \times d}$ and $V \in \mathbb{R}^{n \times d}$ contain the user and item latent factor vectors, and d is the vector dimensionality. Given U and V , the weight matrix and bias vectors of every layer are learned by backpropagation with stochastic gradient descent (SGD). Gradients of W and b are calculated as follows:

$$\frac{\partial L}{\partial W^u} = \lambda_w W^u + \lambda_m \frac{\partial \bar{S}_c}{\partial W^u} \sum_c \text{loss}(S_c^u, \bar{S}_c^u) + \lambda_u \frac{\partial h_{+,0}^u}{\partial W^u} (U - h_{+,0}^u), \quad (7)$$

$$\frac{\partial L}{\partial b^u} = \lambda_w b^u + \lambda_m \frac{\partial \bar{S}_c}{\partial b^u} \sum_c \text{loss}(S_c^u, \bar{S}_c^u) + \lambda_u \frac{\partial h_{+,0}^u}{\partial b^u} (U - h_{+,0}^u). \quad (8)$$

A learning rate α is adopted to update all parameters using calculated gradients.

IV. EXPERIMENTS

Experiments are conducted on three real world datasets, MovieLens-100k (*ml-100k*), MovieLens-10M (*ml-10m*), and one dataset from an e-commerce company (OfflinePay). We first investigate whether the flexible autoencoder architecture of our model can generate more accurate latent representations on non-sequential data. Experiments on OfflinePay evaluate the effectiveness of sequential data modeling.

A. Datasets and preprocessing

The first dataset, *ml-100k*, contains ratings from 943 users on 1,682 movies. It has demographic data for users and descriptions for movies. The second dataset, *ml-10m*, contains 10,000,054 ratings and 95,580 tags from 71,567 users for 10,681 movies. It contains item content information, but no demographic data. We employ user-added tags as an information source for users as well as for movies.

OfflinePay is a dataset of user purchases in (offline) shops, paying with a plastic e-money card. The dataset contains a total of 67M transaction records from a four-month period. The goal of using the OfflinePay dataset is to recommend new shop genres to users, not individual products. After aggregating all transaction data into the format of (user i , shop genre j , number of transactions r_{ij}), and removing shoppers who used only one shop genre, the number of r_{ij} values is 7,150,833 with 961,992 unique users and 105 shop genres. The auxiliary data sources include user registered information and shop genre textual descriptions. Additionally, we collect user purchase history on an e-commerce platform during the same time period. The sequence data contains the genres of purchased items online.

The datasets are preprocessed to fixed-length embeddings for non-sequential data, and sequences of embedding vectors for sequential data, respectively. For *ml-100k*, we discretize continuous features like age to discrete values, compute a bag-of-words vector for each user and item. The vector dimensions are 821 for users, 2,482 for movies, respectively. For *ml-10m*, movie content description and tags that users give to items are textual information. We first tokenize texts, then train Doc2vec vectors for every data source with the embedding vector length set to 500.

To generate shop genre embedding vectors for the OfflinePay dataset, all shop names that belong to same genre are grouped together and Doc2vec is applied to generate a 300-dimensional vector for each shop genre. User registered information is preprocessed the same way as *ml-100k*, and the vector length is 189. For the sequence of genre purchase history, Word2vec is adopted to build 100-d embedding vectors after tokenization. Genres in each sequence are mapped to the corresponding embedding vectors.

In experiments, we rank predicted ratings of candidate items and recommend the top M to each user. Mean average precision (MAP) and recall are used as evaluation metrics.

Fig. 3: **Recall@M comparison on *ml-100k*, *ml-10m*.** Results are shown for *I-AutoRec*, *CDL*, *DCF*, *aSDAE* and *DHA*. The dimension d of the latent factor vector is set to 50, 100 and 150, respectively.

B. Experimental setting

The number of hidden layers of each model is optimized on a validation dataset. The first fusion hidden layer of *DHA* is used to bridge the joint training between feature space learning and collaborative filtering. For other models, if the total number of hidden layers is L , we connect layer $L/2$ for joint training. The number of units in each hidden layer is incremented by K from the middle of the autoencoder to both sides. For sequential data modeling, recent T purchases is used in the experiments, and values $T \in \{5, 10\}$ are evaluated in our experiments.

The mini-batch size is set to 50 and 1,000 for *ml-100k* and *ml-10m*, respectively. For the *OfflinePay* dataset, since the numbers of unique users and items differ significantly, it is set to 20 for items and 10,000 for users, separately. The model is implemented using the Theano library.

C. Experiments on MovieLens datasets

We compare our model with the following algorithms. Note that experiments on *MovieLens* do not include sequential data.

- *AutoRec* [17]: *I-AutoRec* takes a partial item feedback vector as input and reconstructs at the output layer.
- *CDL* [7]: a hierarchical Bayesian model that jointly performs deep representation learning for content information and collaborative filtering for the ratings matrix.
- *DCF* [12]: a model that combines matrix factorization with marginalized denoising stacked autoencoders. We concatenate side information as input to *DCF*.
- *aSDAE* [19]: a hybrid model that integrates side information by an additional denoising autoencoder into the matrix factorization model.

- *DHA*: the proposed model that applies independent autoencoder architecture to heterogeneous data sources.

To compare different models, we repeat 80-20 splits of the data 5 times, run 5-fold cross validation and report average performance. Grid search is applied to find optimal hyperparameters for all models. We search the learning rate of SGD, $\alpha \in \{0.1, 0.05, 0.01, 0.001\}$, the regularization of learned parameters, λ_f and λ_w of our model $\in \{2.0, 0.1, 0.01, 0.001\}$, the corruption level of masking noise $\in \{0.1, 0.3\}$, the activation function $\in \{\text{sigmoid}, \text{relu}\}$, and the number of fusion hidden layers $\in \{1, 2\}$. The parameters used to balance loss between user and item, $\lambda_m, \lambda_n, \lambda_u, \lambda_v$ are set to 1. For *CDL*, *DCF* and *aSDAE*, we search hidden layer number from 4 and 6. The joint training is alternated 5 times, and we run 5 epochs for learning features in each alteration. Before the joint training, layer-wise pretraining is conducted to initialize network weights.

For the experiment on *ml-100k*, we input rating vectors, item content information and user demographic data to *DHA* and *aSDAE*. Rating vectors are not used in *DCF* and only item content information is used in *CDL*. *I-AutoRec* leverages no side information. After grid search, the adopted hidden layer number of *CDL*, *DCF* and *aSDAE* is 4. The number of fusion hidden layer is set to 1 for *DHA*. The parameter for regularizing learned parameters is set to 0.01 in *DHA*, 0.001 in *CDL* and *aSDAE*, and 0.1 in *DCF*, respectively. The optimal performance is found when the learning rate is set to 0.001 for *CDL*, *DCF*, 0.01 for *aSDAE* and *DHA*, and 0.1 for *I-AutoRec*.

As shown in Fig. 3, all models achieve better recall than *I-AutoRec*, showing the advantage of using side information. *DHA* and *aSDAE* perform better than *CDL* which only incor-

TABLE I: **MAP@100 comparison on *ml-100k* and *ml-10m* datasets.** Results are shown for three different settings of user and item latent factor vectors, $d=50, 100$, and 150 .

Model	<i>ml-100k</i>			<i>ml-10m</i>		
	d=50	d=100	d=150	d=50	d=100	d=150
I-AutoRec	0.0573	0.0568	0.0572	0.0325	0.0323	0.0326
CDL	0.1896	0.1825	0.1685	0.1458	0.1532	0.1612
DCF	0.2012	0.2028	0.2069	0.1591	0.1620	0.1566
aSDAE	0.2161	0.2228	0.2142	0.1602	0.1560	0.1642
DHA	0.2236	0.2304	0.2258	0.1793	0.1774	0.1824

porates item content description. DHA outperforms aSDAE which integrates raw side information at every hidden layer. The MAP comparison in Table I shows our model obtains more precise results for all dimension settings.

There are five sets of available inputs for the experiment on the *ml-10m* dataset. Users and movies have rating and tag vectors, movies also have content vectors. For CDL, DCF and aSDAE, different information vectors are concatenated as input. Our model uses all components, *i.e.* two components for users and three components for movies. In the experiment, the best performance is obtained when the number of hidden layers is set to 4 for CDL, aSDAE and to 6 for DCF. In our model, we use 2 fusion hidden layers and different layer numbers for components. The number of hidden layers, L_c , is set to 4 for users and movie rating vectors and to 2 for tag and content vectors. As shown in Fig. 3, DHA obtains better recall performance compared to other algorithms. aSDAE is competitive and outperforms both DCF and CDL in three dimension settings. The MAP comparison in Table I indicates that in addition to producing recommendation with better recall, our model also achieves better precision results.

D. Experiments on *OfflinePay* dataset

Since the *OfflinePay* dataset involves user online purchase histories, we use the first 3-month data as training data, the following half month's data as validation dataset to find optimal parameters, and data from the remaining half-month as test set. We compare the following algorithms:

- implicit-cf [1]: a matrix factorization model for implicit datasets.
- CDL [7]: a Bayesian model that learns a feature space from item information and jointly trains with CF.
- DCF [12]: a model that incorporates side information by marginalized denoising stacked autoencoders with a matrix factorization model.
- DHA-RNNED-s10: our model that learns a latent representation only from the sequence of online purchases. The number of time steps in each sequence is 10.
- DHA-RNNED-s5: our model with the same modeling process as DHA-RNNED-s10, but using 5 time steps in each sequence.
- DHA-RNNED-item: our model extracts features from sequential online purchases at user side, and from shop genre descriptions at item side.

- DHA-all: the proposed model that leverage non-sequential side information sets and sequential online purchase activities simultaneously. The used time step number of the purchase sequence is 10.

In the experiment, the joint learning is alternated 3 times, and we run 3 epochs for feature extraction every time. The number of hidden layers for CDL, DCF and our model is set to 4, and 1 fusion hidden layer is used in DHA models. For the sequential modeling, we set the hidden units of LSTMs to be the same as the dimension of the user and item latent factor vector. The SGD learning rate and regularization parameters for each model are found by grid search on the validation set. We set the learning rate to 0.1 for implicit-cf and CDL, to 0.001 for DCF and to 0.01 for the other models. The parameter to regularize learned parameters is set to 2.0 for CDL, and 0.1 for DCF and DHA-all. There is no training alteration for implicit-cf, but we run 25 iterations to learn user and item latent factor vectors.

DCF integrates both user registration information and shop genre descriptions, while CDL uses only the latter one. DHA-RNNED-s10 and DHA-RNNED-s5 do not include any side information except user online purchases. DHA-RNNED-item adopts sequential data and shop genre descriptions, and DHA-all utilizes all of the data. Note that since ratings are not used in any models, aSDAE is not applied on *OfflinePay* dataset.

From Fig. 4, we observe that models taking advantage of side information have better recall than the baseline implicit-cf. CDL outperforms DCF which, in fact, uses more information sets. This may be due to the fact that many user registration records have outdated or missing values, making the feature extraction less accurate. Compared to CDL and DCF, the proposed models with sequential data modeling achieve better recall. This is due to the fact that offline shop genres in the dataset are included in the online purchased genres. This also indicates that the latent features is able to be extracted from recent online purchases accurately, and reflect the trends of user interests, then lead to better recommendations for offline products, as well. The MAP comparison in Table II shows that the models involving sequential modeling achieve higher precision. This consistently shows that the modeling of online purchases helps with offline product recommendation.

The recall comparison in Fig. 4 shows that DHA-RNNED-

TABLE II: **MAP@100 comparison on *OfflinePay* dataset.** Results are shown for different dimensions of the latent factor vector, $d=50, 100, 150$.

Models	d=50	d=100	d=150
implicit-cf	0.0155	0.0178	0.0177
CDL	0.0296	0.0336	0.0333
DCF	0.0237	0.0311	0.0306
DHA-RNNED-s10	0.0327	0.0333	0.0339
DHA-RNNED-s5	0.0307	0.0343	0.0367
DHA-RNNED-item	0.0394	0.0402	0.0345
DHA-all	0.0424	0.0403	0.0361

Fig. 4: **Recall@M comparison on OfflinePay dataset.** Results are shown for *implicit-CF*, *CDL*, *DCF*, *DHA-RNNED-s10*, *DHA-RNNED-s5*, *DHA-RNNED-item* and *DHA-all*. The dimension of the latent factor vector is set to 50, 100 and 150, respectively.

s10 and DHA-RNNED-s5 have a similar trend as recommended item M increases. These two models use only the sequence of purchased genres from an online e-commerce platform, but with different time steps in the sequence. It is also shown that DHA-all and DHA-RNNED-item have similar recalls. The difference between these two models is that the latter model does not include user registered data. Linking to the previous observation that CDL outperforms DCF, user data does not significantly contribute to the recommendation results.

In order to compare the effect of purchase recency of the input sequence, we apply DHA-RNNED-s10 and DHA-RNNED-s5 to encode the recent ten and five purchases, respectively. Our hypothesis is that more recent online purchases are more representative of current user interests. Although the difference is not big, the recall and MAP comparisons support our hypothesis. The experiments demonstrate that with the independent autoencoder structure for user and item side information and the modeling of user online activities, our model is able to achieve competitive recall and MAP results.

V. CONCLUSIONS

We proposed a model that incorporates multiple sources of heterogeneous auxiliary information in a consistent way to alleviate the data sparsity problem of recommender systems. It takes static and sequential data as input and captures both the inherent tastes of users as well as the dynamics of user preference. The model uses a flexible autoencoder structure for integrating different data sources leading to significant performance gains.

REFERENCES

- [1] Y. Hu, Y. Koren, and C. Volinsky, "Collaborative filtering for implicit feedback datasets," In Proc. Eighth IEEE ICDM, pages 263–272, 2008.
- [2] A. V. D. Oord, S. Dieleman, and B. Schrauwen, "Deep content-based music recommendation," In Proc. 26th International Conference on NIPS, Vol. 2, pp. 2643–2651, 2013.
- [3] F. Zhang, N. J. Yuan, D. Lian, X. Xie, and W. Y. Ma, "Collaborative knowledge base embedding for recommender systems," In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 353–362, 2016.
- [4] S. Oramas, O. Nieto, M. Sordo, and X. Serra, "A deep multimodal approach for cold-start music recommendation" In Proceedings of the 2nd Workshop on Deep Learning for Recommender Systems, 2017.
- [5] I. Porteous, A. Asuncion, and M. Welling, "Bayesian matrix factorization with side information and dirichlet process mixtures," In Proc. 24th AAAI, pp. 563–568, 2010.
- [6] P. Loyola, C. Liu, and Y. Hirate, "Modeling user session and intent with an attention-based encoder-decoder architecture," In Proc. of 11th ACM Conference on Recommender Systems, pp. 147–151, 2017.
- [7] H. Wang, N. Wang, and D. Y. Yeung, "Collaborative deep learning for recommender systems," In Proc. 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1235–1244, 2015.
- [8] P. Vincent, H. Larochelle, Y. Bengio and P.A. Manzagol, "Extracting and composing robust features with denoising autoencoders," In Proc. Twenty-fifth ICML, pp. 1096–1103, 2008.
- [9] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P. A. Manzagol, "Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion," JMLR, Vol. 11, pp. 3371–3408, 2010.
- [10] W. Wang, R. Arora, K. Livescu, and J. Bilmes, "On deep multi-view representation learning", In Proc. 32nd ICML, Vol. 37, pp. 1083–1092, 2015.
- [11] K. Cho, B. V. M., and C. Gulcehre, "Learning phrase representations using RNN encoderdecoder for statistical machine translation," In Proc. 2014 Conference on Empirical Methods in Natural Language Processing, pp. 1724–1734, 2014.
- [12] S. Li, J. Kawale, and Y. Fu, "Deep collaborative filtering via marginalized denoising auto-encoder," In Proc. 24th ACM CIKM, pp. 811–820, 2015.
- [13] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T. S. Chua, "Neural collaborative filtering," In Proc. 26th International Conference on WWW, pp. 173–182, 2017.
- [14] L. Zheng, V. Noroozi, and P. S. Yu, "Joint deep modeling of users and items using reviews for recommendation," In Proc. Tenth ACM ICWDM, pp. 425–434, 2017.
- [15] C. Y. Wu, A. Ahmed, A. Beutel, A. J. Smola, and H. Jing, "Recurrent recommender networks," In Proc. Tenth ACM ICWDM, pp. 495–503, 2017.
- [16] C. Wang, and D. M. Blei, "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.
- [17] S. Sedhain, A. K. Menon, S. Sanner, and L. Xie, "AutoRec: autoencoders meet collaborative filtering," In Proc. 24th International Conference on WWW, pp. 111–112, 2015.
- [18] S. Zhang, L. Yao, and X. Xu, "AutoSVD++: an efficient hybrid collaborative filtering model via contractive autoencoders," In Proc. 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 957–960, 2017.
- [19] X. Dong, L. Yu, Z. Wu, Y. Sun, L. Yuan, and F. Zhang, "A hybrid collaborative filtering model with deep structure for recommender systems," AAAI, 2017.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, "Deep Learning," MIT Press, 2016.