

CoNet: Collaborative Cross Networks for Cross-Domain Recommendation

Guangneng Hu, Yu Zhang, Qiang Yang

Department of Computer Science and Engineering,
Hong Kong University of Science and Technology

njuhgn@gmail.com, yu.zhang.ust@gmail.com, qyang@cse.ust.hk

Abstract. The cross-domain recommendation technique is an effective way of alleviating the data sparse issue in recommender systems by leveraging the knowledge from relevant domains. Transfer learning is a class of algorithms underlying these techniques. In this paper, we propose a novel transfer learning approach for cross-domain recommendation by using neural networks as the base model. In contrast to the matrix factorization based cross-domain techniques, our method is deep transfer learning, which can learn complex user-item interaction relationships. We assume that hidden layers in two base networks are connected by cross mappings, leading to the collaborative cross networks (CoNet). CoNet enables dual knowledge transfer across domains and is achieved in multi-layer feedforward networks which can be trained efficiently by back-propagation. The proposed model is thoroughly evaluated on two large real-world datasets. It outperforms baselines by relative improvements of 7.84% in NDCG. We demonstrate the necessity of adaptively selecting representations to transfer. Our model can reduce tens of thousands training examples without performance degradation by comparing with non-transfer methods.

1 Introduction

Collaborative filtering (CF) approaches, which model the preference of users on items based on their past interactions such as product ratings, are the corner stone for recommender systems. Matrix factorization (MF) is a class of CF methods which learn user latent factors and item latent factors by factorizing their interaction matrix [28,19]. Neural collaborative filtering is another class of CF methods which use neural networks to learn the complex user-item interaction function [9,4,13]. Neural networks have the ability to learn highly nonlinear function, which is suitable to learn the complex user-item interaction. Both traditional MF and neural CF, however, suffer from the cold-start and data sparse issues.

One effective solution is to transfer the knowledge from relevant domains and cross-domain recommendation techniques address such problems [2,21,33,3]. In real life, a user typically participates several systems to acquire different information services. For example, a user installs applications in an app store as well as reads news from a website. It brings us an opportunity to improve the recommendation performance in the target service (and all services) by learning across domains. Following the above example, we can represent the app installation feedback using a binary matrix where the entries indicate whether a user has installed an app. Similarly, we use another binary matrix to indicate whether a user has read a news article. Typically these two matrices are highly sparse, and it is beneficial to learn them simultaneously. This idea is sharpened into the collective matrix factorization (CMF) [37] approach which jointly factorizes these two matrices by sharing the user latent factors. It combines CF on a target domain and another CF on an auxiliary domain, enabling knowledge transfer [31,45]. CMF, however, is a shallow model and has the difficulty in learning the complex user-item interaction function [9,13]. Its knowledge sharing is only limited in the lower level of user latent factors.

Motivated by benefitting from both knowledge transfer learning and learning interaction function, we propose a novel deep transfer learning approach for cross-domain recommendation using neural networks as the base model. Though neural CF approaches are proposed for single domain recommendation [13,40], there are few related works to study knowledge transfer learning for cross-domain recommendation using neural networks. Instead, neural networks have been used as the base model in natural language

processing [5,43] and computer vision [44,27,8]. We explore how to use a neural network as the base model for each domain and enable the knowledge transfer on the entire network across domains. Then a few questions and challenges are raised: 1) What to transfer/share between these individual networks for each domain? 2) How to transfer/share during the learning of these individual networks for each domain? and 3) How is the performance compared with single domain neural learning and shallow cross-domain models?

This paper aims at proposing a novel deep transfer learning approach by answering these questions under cross-domain recommendation scenario. The usual transfer learning approach is to train a base network and then copy its first several layers to the corresponding first layers of a target network with fine-tuning or parameter frozen [44]. This way of transferring has possibly two weak points. Firstly, the shared-layer assumption is strong in practice as we find that it does not work well on real-world cross-domain datasets. Secondly, the knowledge transfer happens in one direction, i.e., only from source to target. Instead, we assume that hidden layers in two base networks are connected by dual mappings, which do not require them to be identical. We enable dual knowledge transfer across domains by introducing cross connections from one base network to another and vice versa, letting them benefit from each other. These ideas are sharpened into the proposed collaborative cross networks (CoNet). CoNet is achieved in simple multi-layer feedforward networks by using dual shortcut connections and joint loss functions, which can be trained efficiently by back-propagation.

The paper is organized as follows. We firstly introduce the preliminaries in Section 2, including notations and the base network. In Section 3, we then present an intuitive model to realize the cross-domain recommendation and point out several intrinsic weaknesses which limit its use. We propose a novel deep transfer learning approach for cross-domain recommendation, named collaborative cross networks (CoNet) in Section 4. The core component is the cross connection units which enable knowledge transfer between source and target networks (Sec. 4.1). Its adaptive variant enforces the sparse structure which adaptively controls when to transfer (Sec. 4.3). In Section 5, we experimentally show the benefits of both transfer learning and deep learning for improving the recommendation performance in terms of ranking metrics (Sec. 5.2). We show the necessity of adaptively selecting representations to transfer (Sec. 5.3). We reduce tens of thousands training examples without performance degradation by comparing with non-transfer models (Sec. 5.4), which can be used to save the cost/labor of labelling data. We review related works in Section 6 and conclude the paper in Section 7.

2 Preliminary

We first give notations and describe the problem setting (Sec. 2.1). We then review a multi-layer neural network as the base network for collaborative filtering (Sec. 2.2).

2.1 Notation

We are given two domains, a source domain \mathcal{S} (e.g. the news recommendation) and a target domain \mathcal{T} (e.g. the app recommendation). As a running example, we let the app recommendation be the target domain and the news recommendation be the source domain. The set of users in two domains are shared, denoted by \mathcal{U} (of size $m = |\mathcal{U}|$). Denote the set of items in \mathcal{S} and \mathcal{T} by \mathcal{I}_S and \mathcal{I}_T (of size $n_S = |\mathcal{I}_S|$ and $n_T = |\mathcal{I}_T|$), respectively. Each domain is a problem of collaborative filtering for implicit feedback [31,16]. For the target domain, let a binary matrix $\mathbf{R}_T \in \mathbb{R}^{m \times n_T}$ describe user-app installing interactions, where an entry $r_{ui} \in \{0, 1\}$ is 1 (observed entries) if user u has an interaction with app i and 0 (unobserved) otherwise. Similarly, for the source domain, let another binary matrix $\mathbf{R}_S \in \mathbb{R}^{m \times n_S}$ describe user-news reading interactions, where the entry $r_{uj} \in \{0, 1\}$ is 1 if user u has an interaction with news j and 0 otherwise. Usually the interaction matrix is very sparse since a user only consumed a very small subset of all items.

For the task of item recommendation, each user is only interested in identifying top-N items. The items are ranked by their predicted scores: $\hat{r}_{ui} = f(u, i|\Theta)$, where f is the interaction function and Θ are model parameters. For matrix factorization (MF) techniques, the match function is the fixed dot product: $\hat{r}_{ui} = \mathbf{P}_u^T \mathbf{Q}_i$, and parameters are latent vectors of users and items $\Theta = \{\mathbf{P}, \mathbf{Q}\}$ where $\mathbf{P} \in \mathbb{R}^{m \times d}$, $\mathbf{Q} \in \mathbb{R}^{n \times d}$ and d is the dimension. For neural CF approaches, neural networks are used to parameterize function f and learn it from interactions:

$$f(\mathbf{x}_{ui}|\mathbf{P}, \mathbf{Q}, \theta_f) = \phi_o(\phi_L(\dots(\phi_1(\mathbf{x}_{ui}))\dots)), \quad (1)$$

where the input $\mathbf{x}_{ui} = [\mathbf{P}^T \mathbf{x}_u, \mathbf{Q}^T \mathbf{x}_i]$ is merged from projections of the user and the item, and the projections are based on their one-hot encodings $\mathbf{x}_u \in \{0, 1\}^m$, $\mathbf{x}_i \in \{0, 1\}^n$ and embedding matrices $\mathbf{P} \in \mathbb{R}^{m \times d}$, $\mathbf{Q} \in \mathbb{R}^{n \times d}$. The output and hidden layers are computed by ϕ_o and $\{\phi_l\}$ in a multilayer feedforward neural network (FFNN), and the connection weight matrices and biases are denoted by θ_f .

In our transfer/multitask learning approach for cross-domain recommendation, each domain is modelled by a neural network and these networks are jointly learned to improve the performance through mutual knowledge transfer. We review the base network in the following subsection before introducing the proposed model.

2.2 Base Network

We adopt an FFNN as the base network to parameterize the interaction function (see Eq.(1)). The base network is similar to the Deep model in [6,4] and the MLP model in [13]. The base network, as shown in Figure 2 (the gray part or the blue part), consists of four modules with the information flow from the input (u, i) to the output \hat{r}_{ui} as follows. **Input** : $(u, i) \rightarrow \mathbf{x}_u, \mathbf{x}_i$. This module encodes user-item interaction indices. We adopt the one-hot encoding. It takes user u and item i , and maps them into one-hot encodings $\mathbf{x}_u \in \{0, 1\}^m$ and $\mathbf{x}_i \in \{0, 1\}^n$ where only the element corresponding to that index is 1 and all others are 0. **Embedding** : $\mathbf{x}_u, \mathbf{x}_i \rightarrow \mathbf{x}_{ui}$. This module embeds one-hot encodings into continuous representations via two embedding matrices and then merges them as $\mathbf{x}_{ui} = [\mathbf{P}^T \mathbf{x}_u, \mathbf{Q}^T \mathbf{x}_i]$ to be the input of successive hidden layers. **Hidden layers**: $\mathbf{x}_{ui} \rightsquigarrow \mathbf{z}_{ui}$. This module takes the continuous representations from the embedding module and then transforms, through multi-hop say L , to a final latent representation $\mathbf{z}_{ui} = \phi_L(\dots(\phi_1(\mathbf{x}_{ui})\dots))$. This module consists of multiple hidden layers to learn nonlinear interaction between users and items. **Output** : $\mathbf{z}_{ui} \rightarrow \hat{r}_{ui}$. This module predicts the score \hat{r}_{ui} for the given user-item pair based on the representation \mathbf{z}_{ui} from the last layer of multi-hop module. Since we focus on one-class collaborative filtering, the output is the probability that the input pair is a positive interaction. This can be achieved by a softmax layer: $\hat{r}_{ui} = \phi_o(\mathbf{z}_{ui}) = 1/(1 + \exp(-\mathbf{h}^T \mathbf{z}_{ui}))$, where \mathbf{h} is parameter.

3 Cross-stitch Networks

We first introduce an intuitive model to realize cross-domain recommendation using neural networks, and point out several intrinsic strong assumptions limiting its use, which inspire the design of our model in the next section.

Intuitively, we can use an MLP model [13] on the target domain and use another MLP model on the source domain. To enable knowledge transfer between two domains, we need some ‘‘cross’’ mapping from the source to the target (and vice versa). We adapt cross-stitch units/networks (CSN) [27] for cross-domain recommendation, which are originally proposed for visual recognition tasks (see Fig. 1a).

Given two activation maps a_A and a_B from the l -th layer for two tasks A and B , CSN learn linear combinations \tilde{a}_A, \tilde{a}_B of both the input activations and feed these combinations as input to the successive layers’ filters:

$$\tilde{a}_A^{ij} = \alpha_S a_A^{ij} + \alpha_D a_B^{ij}, \quad \tilde{a}_B^{ij} = \alpha_S a_B^{ij} + \alpha_D a_A^{ij}, \quad (2)$$

where the shared parameter α_D controls information transferred from the other network, α_S controls information from the task-specific network, and (i, j) is the location in the activation map.

Although the cross-stitch unit indeed incorporates knowledge from the source domain (and target domain vice versa), there are several limitations of this simple stitch unit. Firstly, cross-stitch networks cannot process the case that the dimensions of contiguous layers are different. In other words, it assumes that the activations in successive layers are in the *same vector space*. This is not an issue in convolutional networks for computer vision since the activation maps of contiguous layers are in the same space [20]. For collaborative filtering, however, it is not the case in typical multi-layer FFNNs where the architecture follows a tower pattern: the lower layers are wider and higher layers have smaller number of neurons [6,13]. Secondly, it assumes that the representations from other networks are *equally important* with weights being all the same scalar α_D . Some features, however, are more useful and predictive and it should be learned attentively from data [39]. Thirdly, it assumes that the representations from other networks are *all useful* since it transfers activations from every location in

a dense way. The sparse structure, however, plays a key role in general learning paradigm [8]. Instead, our model can be extended to learn the sparse structure on the task relationship matrices which are defined in Eq. (5), with the help of the existing sparsity-induced regularization. As we will see in the experiments (see Table 2 and Figure 3), the sparse structure is necessary for generalization performance.

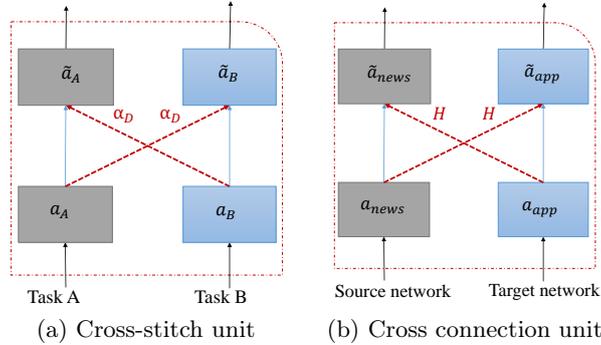


Fig. 1: Left: The cross-stitch unit [27]. Right: The proposed cross connection unit. The shared transfer parameter is a scalar α_D in CSN while it is a matrix \mathbf{H} in CoNet.

4 Collaborative Cross Networks

To alleviate the limitations of cross-stitch networks, we propose collaborative cross networks (CoNet) to transfer knowledge for the cross-domain recommendation. The core component is cross connection units (Sec. 4.1). Our cross unit generalizes the cross-stitch (Sec. 4.2) and exploits the sparse structure (Sec. 4.3). We describe the model learning from implicit feedback data and the optimization process of the joint loss (Sec. 4.4). A complexity analysis is also given (Sec. 4.5).

4.1 Cross Connections Unit

In this section, we present a novel soft-sharing approach for transferring knowledge for cross-domain recommendation. It relaxes the hard-sharing assumption [44] and is motivated by the cross-stitch networks [27].

We now introduce the cross connections unit to enable dual knowledge transfer as shown in Fig. 1b. The central idea is simple, using a matrix rather than a scalar to transfer. Similarly to the cross-stitch network, the target network receives information from the source network and vice versa. In detail, let \mathbf{a}_{app} be the representations of the l -th hidden layer and $\tilde{\mathbf{a}}_{app}$ be the input to the $l + 1$ -th in the app network, respectively. Similarly, they are \mathbf{a}_{news} and $\tilde{\mathbf{a}}_{news}$ in the news network. The cross unit implements as follows:

$$\tilde{\mathbf{a}}_{app} = \mathbf{W}_{app}\mathbf{a}_{app} + \mathbf{H}\mathbf{a}_{news}, \tilde{\mathbf{a}}_{news} = \mathbf{W}_{news}\mathbf{a}_{news} + \mathbf{H}\mathbf{a}_{app}, \quad (3)$$

where \mathbf{W}_{app} and \mathbf{W}_{news} are weight matrices, and the matrix \mathbf{H} controls the information from news network to app network and vice versa. The knowledge transferring happens in two directions, from source to target and from target to source. We enable dual knowledge transfer across domains and let them benefit from each other. When target domain data is sparse, the target network can still learn a good representation from that of the source network through the cross connection units. It only needs to learn “residual” target representations with the reference of source representations, making the target task learning easier and hence alleviating the data sparse issues. The role of matrix \mathbf{H} is similar to the scalar α_D in the sense of enabling knowledge transfer between domains.

We give a closer look at the matrix \mathbf{H} for it can alleviate all three issues faced by cross-stitch unit. Firstly, the successive layers can be in *different vector space* (spaces with different dimensions)

since the matrix \mathbf{H} can be used to match their dimension. For example, if the l -th layer (\mathbf{a}_{app} and \mathbf{a}_{news}) has dimension 128, and the $l + 1$ -th layer ($\tilde{\mathbf{a}}_{app}$ and $\tilde{\mathbf{a}}_{news}$) has dimension 64, then the matrix $\mathbf{H} \in \mathbb{R}^{64 \times 128}$. Secondly, the entries of \mathbf{H} are learned from data. They are likely not to be all the same, showing that the *importances* of transferred representations are different for each neuron/position. Thirdly, we can enforce some prior on the matrix \mathbf{H} to exploit the structure of the neural architecture. The sparse structure can be enforced to adaptively *select useful representations* to transfer. Based on the cross connection units, we propose the CoNet models in the following sections, including a basic model (Sec. 4.2) and an adaptive variant (Sec. 4.3).

4.2 Basic Model

We propose the collaborative cross network (CoNet) model by adding cross connection units (see Sec. 4.1) and the joint loss (see Sec. 4.4) to the entire FFNN, as shown in Figure 2. We firstly describe a basic model in this section and then present an adaptive variant in the next section.

We decompose the model parameters into two parts, task-shared and task-specific: $\Theta_{app} = \{\mathbf{P}, (\mathbf{H}^l)_1^L\} \cup \{\mathbf{Q}_{app}, \theta_{f_{app}}\}$ and $\Theta_{news} = \{\mathbf{P}, (\mathbf{H}^l)_1^L\} \cup \{\mathbf{Q}_{news}, \theta_{f_{news}}\}$ where \mathbf{P} is the user embedding matrix and \mathbf{Q} are the item embedding matrices with the subscript specifying the corresponding domain. We stack the cross connections units on the top of the shared user embeddings, enabling deep knowledge transfer. Denote by \mathbf{W}^l the weight matrix connecting from the l -th to the $l + 1$ -th layer (we ignore biases for simplicity), and by \mathbf{H}^l the linear projection underlying the corresponding cross connections. Then two base networks are coupled by cross connections:

$$\mathbf{a}_{app}^{l+1} = \sigma(\mathbf{W}_{app}^l \mathbf{a}_{app}^l + \mathbf{H}^l \mathbf{a}_{news}^l), \quad (4a)$$

$$\mathbf{a}_{news}^{l+1} = \sigma(\mathbf{W}_{news}^l \mathbf{a}_{news}^l + \mathbf{H}^l \mathbf{a}_{app}^l), \quad (4b)$$

where the function $\sigma(\cdot)$ is the widely used rectified activation units (ReLU) [29]. We can see that \mathbf{a}_{app}^{l+1} receives two information flows: one is from the *transform* gate controlled by \mathbf{W}_{app}^l and one is from the *transfer* gate controlled by \mathbf{H}^l (similarly for the \mathbf{a}_{news}^{l+1} in source network). We call \mathbf{H}^l the relationship/transfer matrix since it learns to control how much sharing is needed. To reduce model parameters and make the model compact, we use the same linear transformation \mathbf{H}^l for two directions, similar to the cross-stitch networks. Actually, using different matrices for two directions did not improve results on the evaluated datasets.

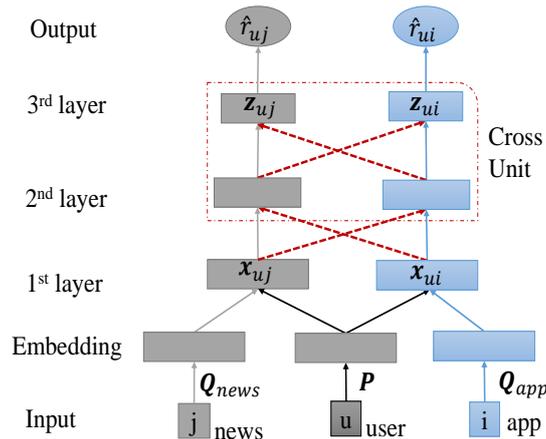


Fig. 2: The proposed collaborative cross networks (a version of three hidden layers and two cross units). We adopt a multilayer FFNN as the base network (grey or blue part, see Sec. 2.2). The red dotted lines indicate the cross connections which enable the dual knowledge transfer across domains (a cross unit is illustrated in the dotted rectangle box, see Fig. ??).

4.3 Adaptive Model

As we can see, the task relationship matrices $\{\mathbf{H}^l\}$ are crucial to the proposed CoNet model. They control the representation transfer from another domain. We can enforce these matrices to have some structure. The assumption is that not all representations from another network are useful. We may expect that the representations coming from other domains are sparse and selective. The selective mechanism can help transfer general/useful representations while ignore the specific/noisy ones. The sparse representations are widely adopted in multitask/transfer learning [1,8,42].

This corresponds to enforcing a sparse prior on the structure and can be achieved by penalizing the task relationship matrix $\{\mathbf{H}^l\}$ via some regularization. It may help the individual network to learn intrinsic representations for itself and other tasks. In other words, $\{\mathbf{H}^l\}$ adaptively controls when to transfer.

We adopt the widely used sparsity-induced regularization—least absolute shrinkage and selection operator (lasso) [38]. In detail, denote by $r \times p$ the size of matrix \mathbf{H}^l (usually $r = p/2$). That is, \mathbf{H}^l linearly transforms representations $\mathbf{a}_{news}^l \in \mathbb{R}^p$ in the news network and the result is as part of the input to the next layer $\tilde{\mathbf{a}}_{app}^{l+1} \in \mathbb{R}^r$ in the app network (see Eq.(4) and Eq.(3)). Denote by h_{ij} the (i, j) entry of \mathbf{H}^l . To induce overall sparsity, we impose the ℓ_1 -norm penalty on the entries $\{h_{ij}\}$ of \mathbf{H}^l :

$$\Omega(\mathbf{H}^l) = \lambda \sum_{i=1}^r \sum_{j=1}^p |h_{ij}|, \quad (5)$$

where hyperparameter λ controls the degree of sparsity. This corresponds to the lasso regularization. We call this sparse variant as the SCoNet model.

Other priors like low-rank factorization are alternatives of sparse structure. And the lasso variants like group lasso and sparse group lasso are also possible. We adopt the general sparse prior and the widely used lasso regularization.

4.4 Model Learning

Due to the nature of the implicit feedback and the task of item recommendation, the squared loss is not suitable since it is usually for rating regression/prediction. Instead, we adopt the cross-entropy loss:

$$\mathcal{L}_0 = - \sum_{(u,i) \in \mathbf{R}^+ \cup \mathbf{R}^-} r_{ui} \log \hat{r}_{ui} + (1 - r_{ui}) \log(1 - \hat{r}_{ui}), \quad (6)$$

where \mathbf{R}^+ and \mathbf{R}^- are the observed interaction matrix and randomly sampled negative examples [31], respectively. This objective function has probabilistic interpretation and is the negative logarithm likelihood of the following likelihood function: $L(\Theta | \mathbf{R}^+ \cup \mathbf{R}^-) = \prod_{(u,i) \in \mathbf{R}^+} \hat{r}_{ui} \prod_{(u,i) \in \mathbf{R}^-} (1 - \hat{r}_{ui})$, where Θ are model parameters.

Now we define the joint loss function, leading to the proposed CoNet model which can be trained efficiently by back-propagation. Instantiating the *base loss* (\mathcal{L}_0) described in Eq. (6) by the *loss of app* (\mathcal{L}_{app}) and *loss of news* (\mathcal{L}_{news}) recommendation, the objective function for the CoNet model is their joint losses: $\mathcal{L}(\Theta) = \mathcal{L}_{app}(\Theta_{app}) + \mathcal{L}_{news}(\Theta_{news})$, where model parameters: $\Theta = \Theta_{app} \cup \Theta_{news}$. Note that Θ_{app} and Θ_{news} share user embeddings and transfer matrices $\{\mathbf{P}, (\mathbf{H}^l)_{l=1}^L\}$. For the CoNet-sparse model, the objective function is added by the term $\Omega(\mathbf{H}^l)$ in Eq.(5).

The objective function can be optimized by stochastic gradient descent and its variants like adaptive moment method (Adam) [17]. The update equations are: $\Theta^{new} \leftarrow \Theta^{old} - \eta \partial L(\Theta) / \partial \Theta$, where η is the learning rate. Typical deep learning library like TensorFlow (<https://www.tensorflow.org>) provides automatic differentiation and hence we omit the gradient equations $\partial L(\Theta) / \partial \Theta$ which can be computed by chain rule in back-propagation.

4.5 Complexity Analysis

The model parameters Θ include $\{\mathbf{P}, (\mathbf{H}^l)_{l=1}^L\} \cup \{\mathbf{Q}_{app}, (\mathbf{W}_{app}^l, \mathbf{b}_{app}^l)_{l=1}^L, \mathbf{h}_{app}\} \cup \{\mathbf{Q}_{news}, (\mathbf{W}_{news}^l, \mathbf{b}_{news}^l)_{l=1}^L, \mathbf{h}_{news}\}$, where the embedding matrices \mathbf{P} , \mathbf{Q}_{app} and \mathbf{Q}_{news} contain a large number of parameters since they depend on the input size of users and items. Typically, the number of neurons in a hidden layer is about one hundred. That is, the size of connection weight matrices and task relationship matrices is

hundreds by hundreds. In total, the size of model parameters is linear with the input size and is close to the size of typical latent factors models [19] and neural CF approaches [13].

During training, we update the target network using the target domain data and update the source network using the source domain data. The learning procedure is similar to the cross-stitch networks [27]. And the cost of learning each base network is approximately equal to that of running a typical neural CF approach [13]. In total, the entire network can be efficiently trained by BP using mini-batch stochastic optimization.

5 Experiment

We conduct thorough experiments to evaluate the proposed models. We show their superior performance over the state-of-the-art recommendation algorithms in a wide range of baselines (Sec. 5.2) and demonstrate the effectiveness of the sparse variant to select representations (Sec. 5.3). We quantify the benefit of knowledge transfer by reducing training examples (Sec. 5.4). Furthermore, we conduct analyses on the sensitivity and sparsity (Sec. 5.5).

5.1 Experimental Setup

We begin the experiments by introducing the datasets, evaluation protocol, baselines, and implementation details.

Dataset We evaluate on two real-world cross-domain datasets. The first dataset, **Mobile**, is provided by a large internet company, i.e., Cheetah Mobile (<http://www.cmcm.com/en-us/>). The information contains logs of user reading news, the history of app installation, and some metadata such as news publisher and user gender collected in one month in the US. The dataset we used contains 1,164,394 user-app installations and 617,146 user-news reading records. There are 23,111 shared users, 14,348 apps, and 29,921 news articles. We aim to improve the app recommendation by transferring knowledge from relevant news reading domain. The data sparsity is over 99.6%. The second dataset is a public **Amazon** dataset (<http://jmcauley.ucsd.edu/data/amazon/>), which has been widely used to evaluate the performance of collaborative filtering approaches [12]. We use the two largest categories, Books and Movies & TV, as the cross-domain. We convert the ratings of 4-5 as positive samples. The dataset we used contains 1,323,101 user-book ratings and 963,373 user-movie ratings. There are 80,763 shared users, 93,799 books, and 35,896 movies. We aim to improve the book recommendation by transferring knowledge from relevant movie watching domain. The data sparsity is over 99.9%. The statistics are summarized in Table 1. As we can see, both datasets are very sparse and hence we hope improve performance by transferring knowledge from auxiliary domains.

Evaluation Protocol For item recommendation task, the leave-one-out (LOO) evaluation is widely used and we follow the protocol in [13]. That is, we reserve one interaction as the test item for each user. We determine hyper-parameters by randomly sampling another interaction per user as the validation/development set. We follow the common strategy which randomly samples 99 (negative) items that are not interacted by the user and then evaluate how well the recommender can rank the test item against these negative ones. Since we aim at top-N item recommendation, the typical evaluation metrics are hit ratio (HR), normalized discounted cumulative gain (NDCG), and mean reciprocal rank (MRR), where the ranked list is cut off at $topN = 10$. HR intuitively measures whether the reserved test item is present on the top-N list, defined as: $HR = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \delta(p_u \leq topN)$, where p_u is the hit position for the test item of user u , and $\delta(\cdot)$ is the indicator function. NDCG and MRR also account for the rank of the hit position, respectively defined as: $NDCG = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{\log 2}{\log(p_u + 1)}$, $MRR = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{p_u}$.

Baseline We compare with various baselines:

Table 1: Datasets and Statistics.

| Dataset | #user | Target Domain | | | Source Domain | | |
|---------|--------|---------------|--------------|---------|---------------|--------------|---------|
| | | #item | #interaction | density | #item | #interaction | density |
| Mobile | 23,111 | 14,348 | 1,164,394 | 0.351% | 29,921 | 617,146 | 0.089% |
| Amazon | 80,763 | 93,799 | 1,323,101 | 0.017% | 35,896 | 963,373 | 0.033% |

| Baselines | Shallow method | Deep method |
|---------------|---------------------|-----------------|
| Single-domain | BPRMF [36] | MLP [13] |
| Cross-domain | CDCF [24], CMF [37] | MLP++, CSN [27] |

Table 2: Comparison results of different methods on two datasets. The best results are boldfaced and the best baselines are marked with stars.

| Dataset | Metric | BPRMF | CMF | CDCF | MLP | MLP++ | CSN | CoNet | SCoNet | improve paired t -test |
|---------|--------|-------|-------|-------|--------|--------|--------|-------|--------------|--------------------------|
| | HR | .6175 | .7879 | .7812 | .8405 | .8445 | .8458* | .8480 | .8583 | 1.47% $p = 0.20$ |
| Mobile | NDCG | .4891 | .5740 | .5875 | .6615 | .6683 | .6733* | .6754 | .6887 | 2.29% $p = 0.25$ |
| | MRR | .4489 | .5067 | .5265 | .6210 | .6268 | .6366* | .6373 | .6475 | 1.71% $p = 0.34$ |
| | HR | .4723 | .3712 | .3685 | .5014 | .5050* | .4962 | .5167 | .5338 | 5.70% $p = 0.02$ |
| Amazon | NDCG | .3016 | .2378 | .2307 | .3143 | .3175* | .3068 | .3261 | .3424 | 7.84% $p = 0.03$ |
| | MRR | .2971 | .1966 | .1884 | .3113* | .3053 | .2964 | .3163 | .3351 | 7.65% $p = 0.05$ |

BPRMF: Bayesian personalized ranking [36] is a typical latent factors CF approach which learns the user and item factors via MF and pairwise rank loss. It is a shallow model and learns on the target domain only. **MLP**: Multilayer perceptron [13] is a typical neural CF approach which learns user-item interaction function using neural networks. MLP corresponds to the base network as described in Section 2.2. It is a deep model and learns on the target domain only. **MLP++**: We combine two MLPs by sharing the user embedding matrix only. This is a degenerated CoNet which has no cross connection units. It is a simple/shallow knowledge transfer approach applied to two domains. **CDCF**: Cross-domain CF with factorization machines (FM) [24] is a state-of-the-art cross-domain recommendation which extends FM [35]. It is a context-aware approach which applies factorization on the merged domains (aligned by the shared users). That is, the auxiliary domain is used as context. On the Mobile dataset, the context for a user in the target app domain is her history of reading news in the source news domain. Similarly, the context for a user in the target book domain is her history of watching movies in the source movie domain on the Amazon dataset. The feature vector for the input is a sparse vector $\mathbf{x} \in \mathbb{R}^{m+n_T+n_S}$ where the non-zero entries are as follows: 1) the index for user id, 2) the index for item id (target domain), and all indices for her reading articles/watching movies (source domain). Since FM can mimic MF, CDCF can be thought of as a single-domain factorization method applied on the merged source-target rating matrix. It showed better performance than other cross-domain methods like triadic (tensor) factorization [15]. It is a shallow cross-domain model. **CMF**: Collective MF [37] is a multi-relation learning approach which jointly factorizes matrices of individual domains. Here, the relation is user-item interaction. On Mobile, the two matrices are $A = \text{“user by app”}$ and $B = \text{“user by news”}$ respectively. The shared user factors P enable knowledge transfer between two domains. Then CMF factorizes matrices A and B simultaneously by sharing the user latent factors: $A \approx P^T Q_A$ and $B \approx P^T Q_B$. It is a shallow model and jointly learns on two domains. This can be thought of a non-deep transfer/multitask learning approach for cross-domain recommendation. **CSN**: The cross-stitch network [27], described in Sect, 3, is a good competitor. It is a deep multitask learning model which jointly learns two base networks. It enables knowledge transfer via a linear combination of activation maps from two networks via a shared coefficient, i.e., α_D in Eq.(2). This is a deep transfer/multitask learning approach for cross-domain recommendation.

Implementation For BPRMF, we use LightFM’s implementation¹ which is a popular CF library. For CDCF, we adapt the official libFM implementation². For MLP, we use the code released by its authors³. For CMF, we use a Python version reference to the original Matlab code⁴. Our methods are implemented using TensorFlow. Parameters are randomly initialized from Gaussian $\mathcal{N}(0, 0.01^2)$. The optimizer is Adam with initial learning rate 0.001. The size of mini batch is 128. The ratio of negative sampling is 1. As for the design of network structure, we adopt a tower pattern, halving the layer size

¹ <https://github.com/lyst/lightfm>

² <http://www.libfm.org>

³ <https://github.com/hexiangnan/neuralcollaborativefiltering>

⁴ <http://www.cs.cmu.edu/ajit/cm/>

for each successive higher layer. Specifically, the configuration of hidden layers in the base network is $[64 \rightarrow 32 \rightarrow 16 \rightarrow 8]$. This is also the network configuration of the MLP model. For CSN, it requires that the number of neurons in each hidden layer is the same. The configuration notation $[64] * 4$ equals $[64 \rightarrow 64 \rightarrow 64 \rightarrow 64]$. We investigate several typical configurations.

5.2 Comparing Different Approaches

In this section, we report the recommendation performance of different methods and discuss the findings. Table 2 shows the results of different models on the two datasets under three ranking metrics. The last two columns are the relative improvement and its paired t -test of our model vs. the best baselines. We can see that our proposed neural models are better than the base network (MLP), the shallow cross-domain models (CMF and CDCF) learned using two domains information, and the deep cross-domain model (MLP++ and CSN) on both datasets.

On Mobile, our model achieves 4.28% improvements in terms of MRR comparing with the non-transfer MLP, showing the benefits of knowledge transfer. Note that, the way of pre-training an MLP on source domain and then transferring user embeddings to target domain as warm-up did not achieve much improvement. In fact, the improvement is so small that it can be ignored. It shows the necessity of dual knowledge transfer in a deep way. Our model improves more than 20% in terms of MRR comparing with CDCF and CMF, showing the effectiveness of deep neural approaches. Together, our neural models consistently give better performance than other existing methods. Within our models (SCoNet vs CoNet), enforcing sparse structure on the task relationship matrices are useful. Note that, the dropout technique and ℓ_2 norm penalty did not achieve these improvements. They may harm the performance in some cases. It shows the necessity of selecting representations.

On Amazon, our model achieves 7.84% improvements in terms of NDCG comparing with the best baselines (MLP++), showing the benefits of knowledge transfer. Compared to the BPRMF, the inferior performance of CMF and CDCF shows the difficulty in transferring knowledge between Amazon Books and Movies, but our models also achieve good results. Comparing MLP++ and MLP, sharing user embedding is slightly better than the base network due to shallow knowledge transfer. Within our models, enforcing sparse structure on the task relationship matrices are also useful.

CSN is inferior to the proposed CoNet models on both datasets. Moreover, it is surprising that the CSN has some difficulty in benefitting from knowledge transfer on the Amazon dataset since it is inferior to the non-transfer base network MLP. The reason is possibly that the assumptions of CSN are not appropriate: all representations from the auxiliary domain are *equally* important and are *all* useful. By using a matrix \mathbf{H} rather than a scalar α_D , we can relax the first assumption. And by enforcing a sparse structure on the matrix, we also relax the second assumption.

Note that the relative improvement of the proposed model vs. the best baseline is more significant on the Amazon dataset than that on the Mobile dataset, though the Amazon is much sparser than the Mobile (see Table 1). One explanation is that the relatedness the book and movie domains is much larger than that between the app and news domains. This will benefit all cross-domain methods including CMF, CDCF, and CSN, since they exploit information from both two domains. Another possibility is that the noise from auxiliary domain proposes a challenge for transferring knowledge. This shows that the proposed model is more effective since it can select useful representations from the source network and ignore the noisy ones. In the next section, we give a closer look at the impact of the sparse structure.

5.3 Impact of Selecting Representations

On two real-world datasets, it both shows the usefulness of enforcing sparse structure on the task relationship matrices \mathbf{H} . We now quantify the contributions of the sparsity to CoNet. We investigate the impact of the sparsity by controlling the difference of architectures between CSN and CoNet. That is, we let them have the same architecture configuration. As a consequence, the performance of ablation comes from different means of knowledge transfer: scalar α_D used in CSN and sparse matrix \mathbf{H} used in SCoNet.

Figure 3 shows the results on the Mobile and Amazon datasets under several typical architectures. We can see that the sparsity contributes to performance improvements and it is necessary to introduce the sparsity in general settings. On the Mobile data, introducing the sparsity improves the NDCG

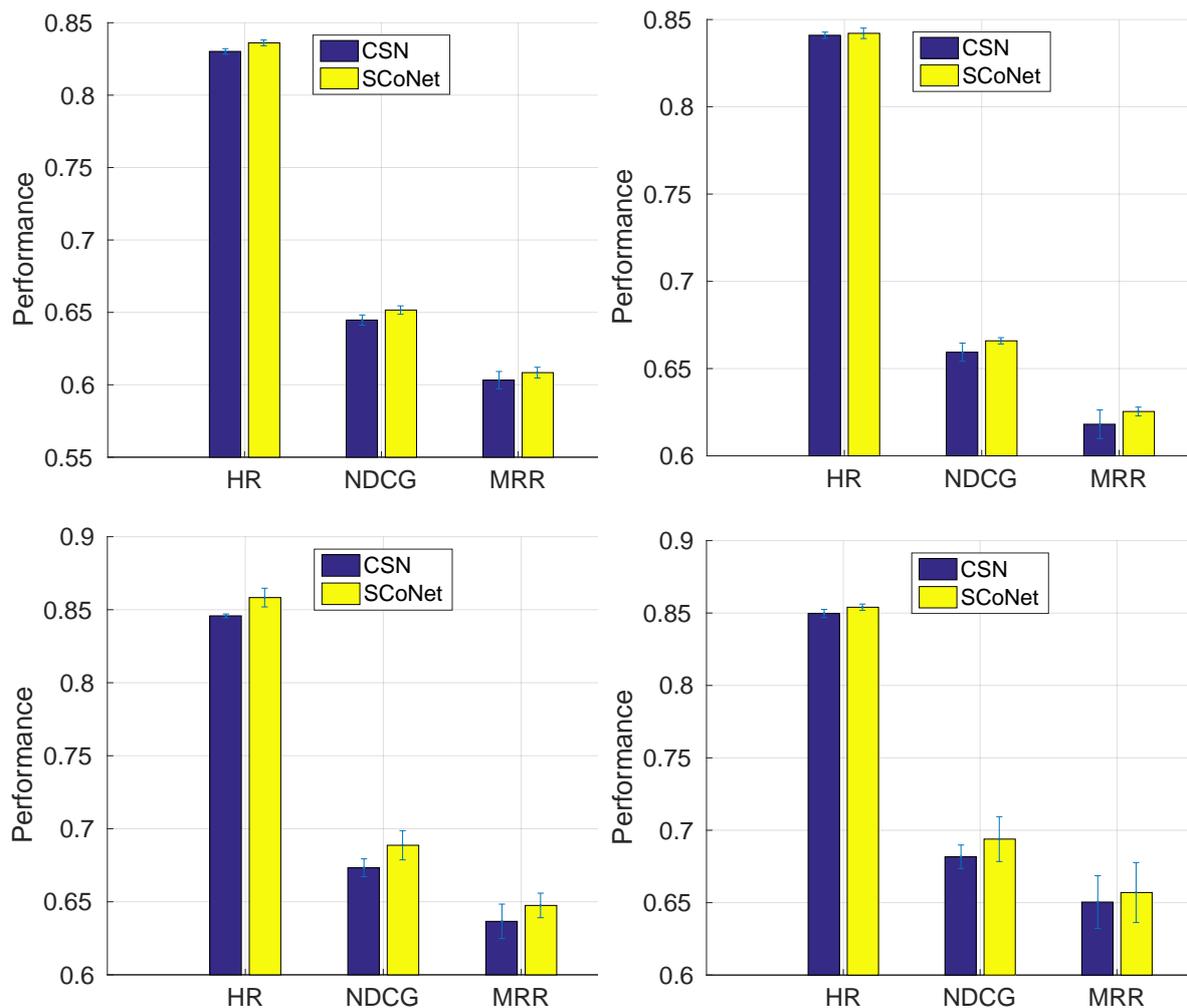


Fig. 3: Impact of the sparsity on the Mobile data. From left to right, configurations are $\{16, 32, 64, 80\} * 4$.

by relatively 2.29%. On the Amazon data, introducing the sparsity improves the NDCG by relatively 4.21%. These results show that it is beneficial to introduce the sparsity and to select representations to transfer on both datasets.

5.4 Benefit of Transferring Knowledge

Transfer learning can reduce the labor and cost of labelling data instances. In this section, we quantify the benefit of knowledge transfer by comparing with non-transfer methods. We do not compare with the cross-domain baselines like CSN in this case because we are to investigate the benefits of transfer learning approaches, not the effectiveness of the proposed model which has demonstrated in the above Sec. 5.2. That is, we gradually reduce the number of training examples in the target domain until the performance of the proposed model is inferior to the non-transfer MLP model. The more training examples we can reduce, the more benefit we can get from transferring knowledge. Note that this is similar to the settings of varying with the cold-start profile size when evaluating new users [18,11].

Referring to Table 1, there are about 50 examples per user on the Mobile dataset. We gradually reduce one and two training examples per user, respectively, to investigate the benefit of knowledge transfer. To be fair, we ensure that every user has at least one training example since the non-transfer MLP cannot deal with the cold-start user issue. The results are shown in Table 3 where the rows corresponding to reduction percentage 0% are copied from Table 2 for clarity. The reduction amount is how many training examples that we remove and the reduction percent is the ratio of reduction amount over the original total training examples. The results show that we can save the cost of labelling about

Table 3: The performance varying with the reduction of training examples. Results with stars are inferior to MLP.

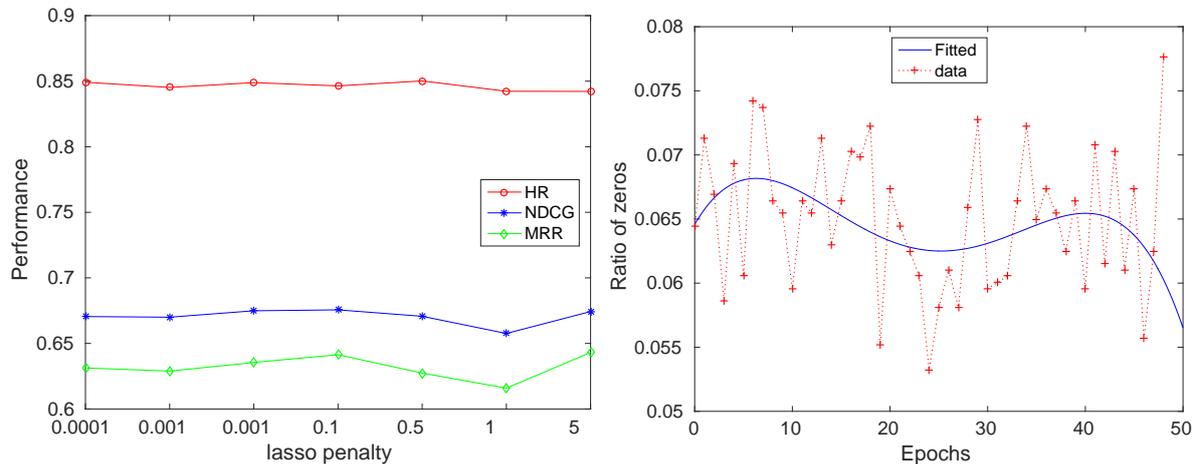
| Dataset | Method | Reduction | | HR | NDCG | MRR |
|---------|--------|-----------|--------|--------|--------|--------|
| | | percent | amount | | | |
| Mobile | MLP | 0% | 0 | .8405 | .6615 | .6210 |
| | | 0% | 0 | .8547 | .6802 | .6431 |
| | SCoNet | 2.05% | 23,031 | .8439 | .6640 | .6238 |
| | | 4.06% | 45,468 | .8347* | .6515* | .6115* |
| Amazon | MLP | 0% | 0 | .5014 | .3143 | .3113 |
| | | 0% | 0 | .5338 | .3424 | .3351 |
| | SCoNet | 1.11% | 12,850 | .5110 | .3209 | .3080* |
| | | 2.18% | 25,318 | .4946* | .3082* | .2968* |

30,000 training examples by transferring knowledge from the news domain but still have comparable performance with the MLP model, a non-transfer baseline.

According to Table 1, there are about 16 examples per user on the Amazon dataset. With a similar setting to the Mobile dataset, the results shown in Table 3 indicates that we can save the cost of labelling about 20,000 training examples by transferring knowledge from movie domain. Note that the Amazon dataset is extremely sparse (the density is only 0.017%), implying that there is difficulty in acquiring many training examples. Under this scenario, our transfer models are an effective way of alleviating the issue of data sparsity and the cost of collecting data.

5.5 Analysis

We analyze the sensitivity to penalty λ in Eq.(5) which controls the sparsity. Results are shown on the Mobile data only due to space limit and we give the corresponding conclusions on the Amazon data. Figure 4 (left) shows the performance varying with the penalty of sparsity enforcing on the task relationship matrices \mathbf{H} . On the Mobile data, the performance achieves good results at 0.1 (default) and 5.0, and it is 0.1 (default) and 1.0 on the Amazon data (not shown).

Fig. 4: Left: Sparsity penalty λ . Right: Ratio of zeros in transfer matrix \mathbf{H} .

Since the sparsity of transfer matrices $(\mathbf{H}^l)_1^L$ is crucial to select representations for transferring, we show the change of zero entries over training epochs. For clarity and due to space limit, we only show the results of the first transfer matrix \mathbf{H}^1 which connects the first and the second hidden layers. Figure 4 (right) shows the results where we use a 4-order polynomial to robustly fit the data. We can see that the matrix becomes sparser for the first 25 iterations, and the general trend is to sparsify.

The average percent of zero entries in \mathbf{H}^1 is 6.5%. For the second and third transfer matrices, the percentage becomes 6.0% and 6.3%, respectively. In summary, sparse transfer matrices are learned and they can adaptively select partial representations to transfer across domains. It may be better to *transfer many instead of all* representations.

6 Related Works

Our work is related to research fields of (cross-domain) recommender systems and (deep) transfer learning.

Recommender systems Recommender systems aim at learning user preferences on unknown items from their past history. Content-based recommendations are based on the matching between user profiles and item descriptions [34]. It is difficult to build the profile for each user when there is no/few content. Collaborative filtering (CF) alleviates this issue by predicting user preferences based on the user-item interaction behavior, agnostic to the content [7,46]. Latent factor models are typical CF methods which learn feature vectors for users and items mainly based on matrix factorization (MF) techniques [19]. MF has probabilistic interpretations [28] and flexible extensions to integrate social relations [25] and item content [26] and both [14], leading to the hybrid methods. Recently, neural networks are proposed to push the learning of feature vectors towards (highly) non-linear representations, learning the user-item interaction function from data rather than using the fixed dot inner used in MF [9,6,13]. Both MF and neural CF models, however, suffer from the data sparsity issue.

Cross-domain recommendation [3] is an effective technique to alleviate sparse issue. A class of methods are based on MF applied to each domain, including the collective MF [37], factorization on both real numbers and binary values matrices [33], and transferring cluster-level patterns transfer [21,22]. Some other shallow methods exploit the bandit learning [23] and the graph-based approach [41]. The deep multi-view neural approach [10] models the shared users as the pivot view and the source/target items as other views by neural networks. We follow this deep learning research thread by using deep networks to learn the interaction function through highly nonlinear transformations.

Transfer and multitask learning Transfer learning (TL) aims at improving the performance of the target domain by exploiting knowledge from source domains [32] which matches the core idea of cross-domain recommendation techniques. The typical TL technique in neural networks is to initialize a target network with transferred features from a pre-trained source network [30,44]. Different from this approach, we transfer knowledge in a deep way such that the two base networks benefit from each other during the learning procedures, motivated by the cross-stitch network [27] which enables information sharing between the two base networks. We generalize it by relaxing the underlying assumption, especially via the idea of selecting representations to transfer.

7 Conclusions

We proposed a novel deep transfer learning for cross-domain recommendation. The sparse target user-item interaction matrix can be reconstructed with the knowledge guidance from the source domain. We demonstrated the necessity of selecting representations to transfer since it may harm the performance by transferring all of them with equal importance. We found that naive deep transfer models may be inferior to the shallow/neural non-transfer methods in some cases. Our transfer models can reduce tens of thousands training examples by comparing with non-transfer methods without performance degradation. Experiments validated their effectiveness by comparing with shallow/deep, single/cross-domain baselines.

The evaluated Mobile dataset is collected from mobile smart devices in different states of the U.S. We found that some hot states like TX, FL, CA, IL and NY have a lot of records while other states are scarce, showing a challenge in reliably learning personalization model in these sparse states. We hypothesize that we can transfer the knowledge from hot states to sparse states, i.e. enabling knowledge transfer between two states. A possible solution is to exploit shared items (apps/news) as a bridge between states, indicating that users in two states are similar if they install/read the same apps/news. The proposed models are then applicable.

Acknowledgment The work is supported by HKPFS PF15-16701.

References

1. A. Argyriou, T. Evgeniou, and M. Pontil. Multi-task feature learning. In *NIPS*, 2007. 4.3
2. S. Berkovsky, T. Kufflik, and F. Ricci. Cross-domain mediation in collaborative filtering. In *UMAP*, 2007. 1
3. I. Cantador, I. Fernández-Tobías, S. Berkovsky, and P. Cremonesi. Cross-domain recommender systems. In *Recommender Systems Handbook*. 2015. 1, 6
4. H.-T. Cheng, L. Koc, J. Harmsen, et al. Wide & deep learning for recommender systems. In *Workshop on DL for RecSys*, 2016. 1, 2.2
5. R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, 2008. 1
6. P. Covington, J. Adams, and E. Sargin. Deep neural networks for youtube recommendations. In *ACM RecSys*, 2016. 2.2, 3, 6
7. M. Deshpande and G. Karypis. Item-based top-n recommendation algorithms. 2004. 6
8. C. Doersch and A. Zisserman. Multi-task self-supervised visual learning. In *ICCV*, 2017. 1, 3, 4.3
9. G. Dziugaite and D. Roy. Neural network matrix factorization. 2015. 1, 6
10. A. Elkahky, Y. Song, and X. He. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *WWW*, 2015. 6
11. I. Fernández-Tobías, M. Braunhofer, M. Elahi, F. Ricci, and I. Cantador. Alleviating the new user problem in collaborative filtering by exploiting personality information. *UMUAI*, 2016. 5.4
12. R. He and J. McAuley. Vbpr: visual bayesian personalized ranking from implicit feedback. In *AAAI*, 2016. 5.1
13. X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *WWW*, 2017. 1, 2.2, 3, 3, 4.5, 5.1, 5.1, 6
14. G. Hu, X. Dai, Y. Song, S. Huang, and J. Chen. A synthetic approach for recommendation: Combining ratings, social relations, and reviews. In *IJCAI*, 2015. 6
15. L. Hu, J. Cao, G. Xu, L. Cao, Z. Gu, and C. Zhu. Personalized recommendation via cross-domain triadic factorization. In *WWW*, 2013. 5.1
16. Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *IEEE ICDM*, 2008. 2.1
17. D. Kingma and J. Ba. Adam: A method for stochastic optimization. 2015. 4.4
18. D. Kluver and J. Konstan. Evaluating recommender behavior for new users. In *ACM RecSys*, 2014. 5.4
19. Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 2009. 1, 4.5, 6
20. A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 3
21. B. Li, Q. Yang, and X. Xue. Can movies and books collaborate?: cross-domain collaborative filtering for sparsity reduction. In *IJCAI*, 2009. 1, 6
22. B. Li, X. Zhu, R. Li, C. Zhang, X. Xue, and X. Wu. Cross-domain collaborative filtering over time. In *IJCAI*, 2011. 6
23. B. Liu, Y. Wei, Y. Zhang, Z. Yan, and Q. Yang. Transferable contextual bandit for cross-domain recommendation. In *AAAI*, 2018. 6
24. B. Loni, Y. Shi, M. Larson, and A. Hanjalic. Cross-domain collaborative filtering with factorization machines. In *ECIR*, 2014. 5.1, 5.1
25. H. Ma, H. Yang, M. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *CIKM*, 2008. 6
26. J. McAuley and J. Leskovec. Hidden factors and hidden topics: understanding rating dimensions with review text. In *ACM RecSys*, 2013. 6
27. I. Misra, A. Shrivastava, A. Gupta, and M. Hebert. Cross-stitch networks for multi-task learning. In *CVPR*, 2016. 1, 3, 1, 4.1, 4.5, 5.1, 5.1, 6
28. A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *NIPS*, 2008. 1, 6
29. V. Nair and G. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 4.2
30. M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *CVPR*, 2014. 6
31. R. Pan, Y. Zhou, B. Cao, N. Liu, R. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *IEEE ICDM*, 2008. 1, 2.1, 4.4
32. S. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 2010. 6
33. W. Pan, N. Liu, E. Xiang, and Q. Yang. Transfer learning to predict missing ratings via heterogeneous user feedbacks. In *IJCAI*, 2011. 1, 6
34. M. Pazzani and D. Billsus. Content-based recommendation systems. In *The adaptive web*. 2007. 6
35. S. Rendle. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology*, 2012. 5.1

36. S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, 2009. 5.1, 5.1
37. A. Singh and G. Gordon. Relational learning via collective matrix factorization. In *SIGKDD*, 2008. 1, 5.1, 5.1, 6
38. R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B*, 1996. 4.3
39. K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 3
40. C. Yang, L. Bai, C. Zhang, Q. Yuan, and J. Han. Bridging collaborative filtering and semi-supervised learning: A neural approach for poi recommendation. In *SIGKDD*, 2017. 1
41. D. Yang, J. He, H. Qin, Y. Xiao, and W. Wang. A graph-based recommendation across heterogeneous domains. In *CIKM*, 2015. 6
42. Z. Yang, B. Dhingra, K. He, W. Cohen, R. Salakhutdinov, and Y. LeCun. Glomo: Unsupervisedly learned relational graphs as transferable representations. 2018. 4.3
43. Z. Yang, R. Salakhutdinov, and W. Cohen. Transfer learning for sequence tagging with hierarchical recurrent networks. 2017. 1
44. J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In *NIPS*, 2014. 1, 4.1, 6
45. Y. Zhang and Q. Yang. A survey on multi-task learning. 2017. 1
46. Z.-D Zhao and M.-S. Shang. User-based collaborative-filtering recommendation algorithms on hadoop. In *International Conference on Knowledge Discovery and Data Mining*, 2010. 6