

From Zero-Shot Learning to Cold-Start Recommendation

Jingjing Li,¹ Mengmeng Jing,¹ Ke Lu,¹ Lei Zhu,² Yang Yang,¹ Zi Huang³

¹University of Electronic Science and Technology of China

²Shandong Normal University; ³The University of Queensland

lijin117@yeah.net; jingmeng1992@gmail.com; kel@uestc.edu.cn; leizhu0608@gmail.com

Abstract

Zero-shot learning (ZSL) and cold-start recommendation (CSR) are two challenging problems in computer vision and recommender system, respectively. In general, they are independently investigated in different communities. This paper, however, reveals that ZSL and CSR are two extensions of the same intension. Both of them, for instance, attempt to predict unseen classes and involve two spaces, one for direct feature representation and the other for supplementary description. Yet there is no existing approach which addresses CSR from the ZSL perspective. This work, for the first time, formulates CSR as a ZSL problem, and a tailor-made ZSL method is proposed to handle CSR. Specifically, we propose a Low-rank Linear Auto-Encoder (LLAE), which challenges three cruxes, i.e., domain shift, spurious correlations and computing efficiency, in this paper. LLAE consists of two parts, a low-rank encoder maps user behavior into user attributes and a symmetric decoder reconstructs user behavior from user attributes. Extensive experiments on both ZSL and CSR tasks verify that the proposed method is a win-win formulation, i.e., not only can CSR be handled by ZSL models with a significant performance improvement compared with several conventional state-of-the-art methods, but the consideration of CSR can benefit ZSL as well.

Introduction

From a near-infinity inventory, recommender systems (Bobadilla et al. 2013; Li et al. 2017) pick a fraction of items which a user might enjoy based on the user’s current context and past behavior (Smith and Linden 2017). If the past behavior, however, is not available, e.g., for a new user, most recommender systems, especially those popular ones based on collaborative filtering (CF) (Ekstrand et al. 2011), would be stuck. Different solutions have been proposed to handle this problem, which is widely known as cold-start recommendation (CSR) (Lin et al. 2013). Recently, cross-domain information (Fernández-Tobías et al. 2012), personal information (Fernández-Tobías et al. 2016) and social network data (Sedhain et al. 2017) have been used to facilitate CSR.

If we take a close look at previous work, it is not hard to find out that the very basic idea behind existing CSR meth-

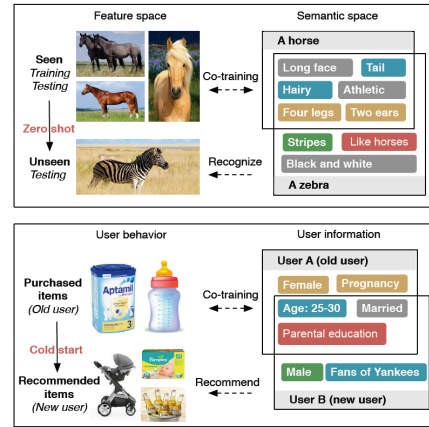


Figure 1: Illustration of ZSL and CSR. It can be seen that they are two extensions of the same intension. However, there is no existing work which addresses CSR from the ZSL perspective. For the first time, CSR is formulated as a ZSL problem and a tailor-made method is proposed to solve CSR in this paper.

ods (Lin et al. 2013; Li et al. 2017) is to *leverage user preferences to generate recommendations for new users*. This is quite reasonable, simply because you cannot deliver the *right* thing to a person you barely know. With the user preferences, we will then have two spaces, an attribute (e.g., user preferences and personal information) space and a behavior (e.g., purchase behavior and past interactions) space, in the cold-start recommendation. The attribute space is used to describe the user preferences, and the behavior space is used to represent the user interactions in the target system. As a result, cold-start recommendation can be defined as a problem to generate recommendations for a fresh user where we have nothing about the user in the behavior space but some side information about the user in the attribute space. With the assumption that *people with the similar preferences would have the similar consuming behavior*, cold-start recommendation can be done by two steps: 1) mapping the behavior space to the attribute space, so that the new users can be linked with the old users; 2) reconstructing user behavior by user attributes, so that we can generate recommendations

for new users. Does it ring a bell now? It is a special case of zero-shot learning (ZSL) (Kodirov, Xiang, and Gong 2017; Ding, Shao, and Fu 2017; Yang et al. 2016).

This paper, for the best of our knowledge, is the first one to investigate CSR in light of ZSL. From Fig. 1, we can clearly see that CSR and ZSL are two extensions of the same intension. Specifically, *both of them involve two spaces, one for direct feature representation and the other for supplementary description, and both of them attempt to predict unseen cases in the feature space by leveraging the description space shared by both seen and unseen ones.* However, CSR and ZSL are not being connected ever before. In this paper, we propose a novel ZSL method to handle the CSR problem.

By formulating CSR as a ZSL problem, we challenge three cruxes in this paper. The first one is the domain shift problem. Not only are the behavior space and the attribute space heterogeneous but also the old users and the new users are divergent in probability distribution. Thus, we have to guarantee that the user behavior can be reconstructed by the user attributes. The second crux is that user behavior in CSR, differs from ZSL in computer vision tasks, is incredibly sparse. In real-world retail giants, such as amazon.com, there are hundreds of millions of users and even more items. A specific user, however, only has a small number of interactions in the system with even less items. In consequence, the user-item matrix would be pretty sparse. The last challenge lies in efficiency. Recommender systems are on-line systems, and people hate waiting.

Technically, we propose a novel ZSL model, named as Low-rank Linear AutoEncoder (LLAE), based on the encoder-decoder paradigm (Kodirov, Xiang, and Gong 2017; Boureau et al. 2007) to handle CSR problems. LLAE consists of an encoder which maps the user behavior space into the user attribute space, and a decoder which reconstructs the user behavior by the user attribute. The reconstruction part guarantees that the user behavior can be generated from user attributes, so that the domain shifts between user behavior and user attributes can be mitigated. We formulate LLAE as a linear model for the efficiency, the computational cost of our model is irrelevant to the number of samples. As a result, it can be applied to large-scale datasets. Furthermore, a low-rank constraint is deployed to handle sparse issues. Low-rank representation (Li et al. 2016) has proven to be efficient for the problem of revealing true data from corrupted observations. We know that a behavior can be linked with numerous attributes, while these attributes should have different weights, and some of the attributes are trivial. If we consider all the attributes, it may weaken the dominant factors, introduce over-fitting and relax the generalization ability. The low-rank constraint, for its mathematical property, helps reveal the dominant factors and filter out trivial connections, or in other words, spurious correlations. It is worth noting that low-rank constraint also helps align the domain shifts from the view of domain adaptation (Li et al. 2018a; 2018b). In summary, the contributions of this paper are:

1) We reveal that CSR and ZSL are two extensions of the same intension. CSR, for the first time, is formulated and

solved as a ZSL problem. Our work builds a connection between CSR and ZSL by cross domain transfer, so that the advances in the two communities can be shared. For instance, when someone who focuses on ZSL noticed our work, he might want to look through recent publications on CSR for inspiration, and vice versa.

- 2) A tailor-made ZSL model, low-rank linear autoencoder (LLAE), is presented to handle the challenging CSR tasks. LLAE takes the efficiency into account, so that it suits large-scale problem.
- 3) Extensive experiments on both ZSL and CSR tasks demonstrate the effectiveness of our method. Excitingly, we find that not only can CSR be handled by ZSL models with a significant performance improvement, but the consideration of CSR can benefit ZSL as well. By linking CSR and ZSL, we wish that this work will benefit both of the communities and elicit more contributions.

Related Work

Zero-shot learning. A basic assumption behind conventional visual recognition algorithms is that some instances of the test class are included in the training set, so that other test instances can be recognized by learning from the training samples. For a large-scale dataset, however, collecting training samples for new and rare objects is painful. A curious mind may ask if we can recognize an unseen object with some semantic description just like human beings do. To that end, zero-shot learning (Zhang and Saligrama 2016; Ding, Shao, and Fu 2017) has been proposed. Typically, ZSL algorithms learn a projection which maps visual space to the semantic space, or the reverse. Different models are proposed based on different projection strategies. From a macro perspective, existing ZSL methods can be grouped into three categories: 1) Learning a mapping function from the visual space to the semantic space (Lampert, Nickisch, and Harmeling 2014; Ding, Shao, and Fu 2017); 2) Learning a mapping function from the semantic space to the visual space (Kodirov et al. 2015); 3) Learning a latent space which shared by the visual domain and the semantic domain (Zhang and Saligrama 2015; 2016).

Cold-start recommendation. Among the models which address cold-start recommendation, we focus on the ones which exploit side information, e.g., user attributes, personal information and user social network data, to facilitate the cold-start problem. Those models can be roughly grouped into three categories, e.g., the similarity based models (Sedhain et al. 2014; Rohani et al. 2014), matrix factorization models (Krohn-Grimberghe et al. 2012; Noel et al. 2012) and feature mapping models (Gantner et al. 2010).

Matrix factorization models typically factorize the relationship matrix into two latent representations by optimizing the following objective:

$$\min_{\mathbf{U}, \mathbf{V}} \|\mathbf{Y} - \mathbf{UV}\|_F^2 + \Omega(\mathbf{U}, \mathbf{V}), \quad (1)$$

where Ω is the regularization used to avoid over-fitting. For cold-start problems, one can learn a shared \mathbf{U} or \mathbf{V} from the side-information, and then use it to predict \mathbf{Y} .

Table 1: Notations and corresponding descriptions, where n , d and k denote the number of samples and dimensionality of behavior space and attribute space, r is the rank of a matrix.

| Notation | Description |
|-----------------------------------------------------------|-----------------------------------------------------|
| $\mathbf{X} \in \mathbb{R}^{d \times n}$ | user behavior space (CSR) / visual space (ZSL) |
| $\mathbf{S} \in \mathbb{R}^{k \times n}$ | user attribute space (CSR) / semantic space (ZSL) |
| $\mathbf{W} \in \mathbb{R}^{k \times d}$ | encoder |
| $\mathbf{M}, \mathbf{W}^\top \in \mathbb{R}^{d \times k}$ | decoder |
| $\mathbf{V} \in \mathbb{R}^{k \times (k-r)}$ | the singular vectors of $\mathbf{W}\mathbf{W}^\top$ |
| $\lambda > 0, \beta > 0$ | penalty parameters |

Feature mapping models normally learn a feature mapping between the side-information and one of the latent feature representations. The differences between the matrix factorization models and the feature mapping models is that in matrix factorization models the shared \mathbf{U} is jointly learned from \mathbf{Y} and the side-information, while in feature mapping models, one needs to learn an additional feature mapping, and further learn different \mathbf{U}_s and \mathbf{U}_t by sharing the feature mapping. More details can be found in (Sedhain et al. 2017; Gantner et al. 2010).

Problem Formulation

Notations

In this paper, we use bold lowercase letters to represent vectors, bold uppercase letters to represent matrices. For a matrix \mathbf{M} , its Frobenius norm is defined as $\|\mathbf{M}\|_F = \sqrt{\sum_i \delta_i(\mathbf{M})^2}$, where $\delta_i(\mathbf{M})$ is the i -th singular value of the matrix \mathbf{M} . The trace of matrix \mathbf{M} is represented by the expression $\text{tr}(\mathbf{M})$. For clarity, we also show the frequently used notations in Table 1.

Linear Low-rank Denoising Autoencoder

Given an input data matrix \mathbf{X} , suppose that we can learn a mapping \mathbf{W} which projects matrix \mathbf{X} into a latent space \mathbf{S} , and another mapping \mathbf{M} which can reconstructs \mathbf{X} from \mathbf{S} . As an optimization problem, our aim is to minimize the reconstruction error. Thus, we have the following objective:

$$\min_{\mathbf{W}, \mathbf{M}} \|\mathbf{X} - \mathbf{M}\mathbf{W}\mathbf{X}\|_F^2, \quad s.t. \mathbf{W}\mathbf{X} = \mathbf{S}. \quad (2)$$

Generally, the latent space is represented as hidden layers. For the concern of efficiency and interpretability, we only deploy one hidden layer \mathbf{S} in our model. In this paper, \mathbf{S} has definite meanings, semantic space in ZSL or user side information in CSR. Recently, tied weights (Mohamed, Dahl, and Hinton 2012) has been introduced into autoencoders to learn faster models yet with less parameters. In this paper, we consider the tied weights $\mathbf{M} = \mathbf{W}^\top$. Then, we have the following formulation as illustrated in Fig. 2:

$$\min_{\mathbf{W}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{W}\mathbf{X}\|_F^2, \quad s.t. \mathbf{W}\mathbf{X} = \mathbf{S}. \quad (3)$$

As we stated before, one of the challenge problems in real-world recommender system is that we need to handle very high-dimensional and sparse matrix, because there are millions of items and users but a specific user only have few

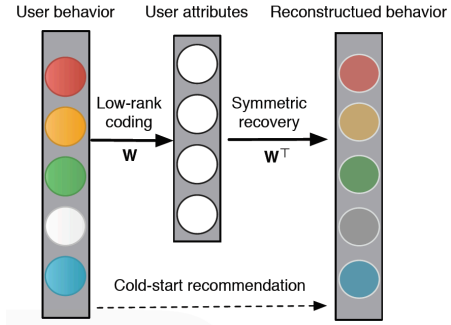


Figure 2: Illustration of the proposed LLAE. Firstly, we learn a low-rank encoder which maps user behavior into user attributes. Then, user attributes of new users are used to reconstruct the user behavior (generate recommendations for new users). For the sake of efficiency, the parameters of the encoder and decoder are symmetric (tied weights). Notice that the encoder guarantees that warm users and cold users can be compared in the attribute space. The reconstruction, at the same time, assures that the user behavior (recommendation list) can be generated from user attributes.

interactions with few items. To avoid spurious correlations caused by the mapping matrix, we propose that \mathbf{W} should be low-rank. As a result, we have:

$$\min_{\mathbf{W}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{W}\mathbf{X}\|_F^2 + \beta \text{rank}(\mathbf{W}), \quad s.t. \mathbf{W}\mathbf{X} = \mathbf{S}, \quad (4)$$

where $\text{rank}(\cdot)$ is the rank operator of a matrix, $\beta > 0$ is a penalty parameter. It is worth noting that the rank constraint on \mathbf{W} benefits the model from at least two aspects. On one side, it helps filter out the spurious correlations from behavior space to attribute space. On the other side, it helps highlight the shared attributes across different users. For instance, a specific attribute, e.g., basketball fan, would be shared by many users from different ages. The low-rank constraint on \mathbf{W} will reveal these common attributes.

In some cold-start tasks, the two spaces might be not very correlated. The low-rank constraint helps reveal the most correlated (relatively) part (dominate factors), and the reconstruction part is even more critical because the projection can be spurious in this situation without reconstruction constraint. The reconstruction part is effective in mitigating the domain shift problem. This is because although the user behavior may change from warm users to cold users, the demand for more truthful reconstruction from the attributes to behavior is generalizable across warm and cold domains, resulting in the learned project function less susceptible to domain shift.

The low-rank constraint on \mathbf{W} makes the optimization more difficult for the reason that low-rank is a well-known NP-hard problem. As an alternative method, the trace norm $\|\cdot\|_*$ is widely used to encourage low-rankness in previous work (Li et al. 2016). However, the trace norm controls the single values of the matrix, but the changes of the single values do not always lead to a change of the rank. Inspired by

recent work (Ding, Shao, and Fu 2017), we propose to use an explicit form of low-rank constraint as follows:

$$\begin{aligned} \min_{\mathbf{W}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{W} \mathbf{X}\|_F^2 + \beta \sum_{i=r+1}^d (\sigma_i(\mathbf{W}))^2, \\ \text{s.t. } \mathbf{W} \mathbf{X} = \mathbf{S}, \end{aligned} \quad (5)$$

where $\sigma_i(\mathbf{W})$ is the i -th singular value of \mathbf{W} , d is the total number of singular values of \mathbf{W} . Different from the trace norm, $\sum_{i=r+1}^d (\sigma_i(\mathbf{W}))^2$ explicitly solves the problem of minimizing the square sum of r -smallest singular value of the mapping \mathbf{W} .

Note that

$$\sum_{i=r+1}^d (\sigma_i(\mathbf{W}))^2 = \text{tr}(\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}), \quad (6)$$

where $\text{tr}(\cdot)$ is the trace operator of a matrix, and \mathbf{V} consists of the singular vectors which correspond to the $(d-r)$ -smallest singular values of $\mathbf{W} \mathbf{W}^\top$. Thus, our objective function can be written as:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{V}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{W} \mathbf{X}\|_F^2 + \beta \text{tr}(\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}), \\ \text{s.t. } \mathbf{W} \mathbf{X} = \mathbf{S}. \end{aligned} \quad (7)$$

At last, to learn more robust hidden layers, we train a denoising autoencoder (Vincent et al. 2008) by introducing corruptions into the input. Specifically, we randomly set 10% of \mathbf{X} to zeros, and get the corrupted version $\hat{\mathbf{X}}$. As a result, we have the final objective:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{V}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{W} \mathbf{X}\|_F^2 + \beta \text{tr}(\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}), \\ \text{s.t. } \mathbf{W} \hat{\mathbf{X}} = \mathbf{S}. \end{aligned} \quad (8)$$

Problem Optimization

To optimize Eq. (8), we first rewrite it to the following equivalent form:

$$\min_{\mathbf{W}, \mathbf{V}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{S}\|_F^2 + \beta \text{tr}(\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}), \text{ s.t. } \mathbf{W} \hat{\mathbf{X}} = \mathbf{S}. \quad (9)$$

However, the constraint on Eq. (9) is hard to optimize. Here we relax the constraint and get the following optimization problem:

$$\min_{\mathbf{W}, \mathbf{V}} \|\mathbf{X} - \mathbf{W}^\top \mathbf{S}\|_F^2 + \lambda \|\mathbf{W} \hat{\mathbf{X}} - \mathbf{S}\|_F^2 + \beta \text{tr}(\mathbf{V}^\top \mathbf{W} \mathbf{W}^\top \mathbf{V}), \quad (10)$$

where $\lambda > 0$ is a penalty parameter. As a result, Eq. (10) is a convex problem which has global optimal solution. Then, we calculate the derivative of Eq. (10) with respect to \mathbf{W} and set it to zero,

$$\begin{aligned} -\mathbf{S}(\mathbf{X}^\top - \mathbf{S}^\top \mathbf{W}) + \lambda(\mathbf{W} \hat{\mathbf{X}} - \mathbf{S})\hat{\mathbf{X}}^\top + \beta \mathbf{V} \mathbf{V}^\top \mathbf{W} = \mathbf{0}, \\ \Rightarrow (\mathbf{S} \mathbf{S}^\top + \beta \mathbf{V} \mathbf{V}^\top) \mathbf{W} + \lambda \mathbf{W} \hat{\mathbf{X}} \hat{\mathbf{X}}^\top = \mathbf{S}(\mathbf{X}^\top + \lambda \hat{\mathbf{X}}^\top) \end{aligned} \quad (11)$$

It is worth noting that the optimization of \mathbf{W} involves \mathbf{V} . As an optimization trick (Ding, Shao, and Fu 2017), we choose to alternatively update them. At first, by treating \mathbf{V} as a constant, we calculate the derivative w.r.t \mathbf{W} and set it to zero, as shown in Eq. (11). Then, we update \mathbf{V} by Eq. (6).

At last, if we use the following substitutions:

$$\begin{cases} \mathbf{A} = \mathbf{S} \mathbf{S}^\top + \beta \mathbf{V} \mathbf{V}^\top \\ \mathbf{B} = \lambda \mathbf{X} \hat{\mathbf{X}}^\top \\ \mathbf{C} = \mathbf{S} \mathbf{X}^\top + \lambda \mathbf{S} \hat{\mathbf{X}}^\top \end{cases}, \quad (12)$$

Algorithm 1. Low-rank Linear AutoEncoder for CSR

Input: User behavior space \mathbf{X} , user attribute space \mathbf{S} , parameters λ and β .

Output: Recommended items for new users.

Warm-up:

Repeat

1. Solve the eigen-decomposition problem to get \mathbf{V} :
 $\mathbf{V} \leftarrow \text{svd}(\mathbf{W} \mathbf{W}^\top)$.

2. Optimize the encoder \mathbf{W} (and the decoder \mathbf{W}^\top):
 $\mathbf{W} = \text{sylvester}(\mathbf{A}, \mathbf{B}, \mathbf{C})$,

where $\mathbf{A} = \mathbf{S} \mathbf{S}^\top + \beta \mathbf{V} \mathbf{V}^\top$, $\mathbf{B} = \lambda \mathbf{X} \hat{\mathbf{X}}^\top$,
 $\mathbf{C} = \mathbf{S} \mathbf{X}^\top + \lambda \mathbf{S} \hat{\mathbf{X}}^\top$.

Until Convergence

Cold-start:

$\mathbf{X}_{new} = \mathbf{W}^\top \mathbf{S}_{new}$.

Recommendation:

Using the logistic regression function to predict the recommendation probability of items, and recommend the top- k items.

Eq. (11) can be written as the following equation:

$$\mathbf{A} \mathbf{W} + \mathbf{W} \mathbf{B} = \mathbf{C}, \quad (13)$$

which can be effectively solved by Sylvester¹ operation in Matlab with only one line of code,

$$\mathbf{W} = \text{sylvester}(\mathbf{A}, \mathbf{B}, \mathbf{C}).$$

For clarity, we show the proposed method in Algorithm 1.

Complexity Analysis

The computational cost of our algorithm consists of two parts: 1) the optimization of \mathbf{W} and 2) the updating of \mathbf{V} . Generally, both of them cost $O(d^3)$. However, if we directly calculate $\mathbf{V} \mathbf{V}^\top$ instead of \mathbf{V} , the cost of 2) can be reduced to $O(r^2 d)$ ($r \ll d$ is the rank of \mathbf{W}) (Ding, Shao, and Fu 2017). In any case, the complexity of our algorithm only depends on the dimensionality. It is irrelevant to the number of samples. As a result, it can be applied to large-scale datasets.

Zero-shot Classification

Given a training data \mathbf{X} and a semantic representation \mathbf{S} , we can learn an encoder \mathbf{W} and a decoder \mathbf{W}^\top by Eq. (13). For the new test sample set \mathbf{X}_{new} , we can embed it to a semantic space $\mathbf{S}_{new} = \mathbf{W} \mathbf{X}_{new}$. Then, the labels of \mathbf{X}_{new} can be learned by a classifier which calculates the distances between \mathbf{S}_{new} and \mathbf{S}_{proto} , where \mathbf{S}_{proto} is the projected prototypes in the semantic space.

$$f(\mathbf{X}_{new}) = \arg \min \text{dis}(\mathbf{S}_{new}, \mathbf{S}_{proto}), \quad (14)$$

where $f(\mathbf{X}_{new})$ is a classifier which returns the labels of \mathbf{X}_{new} , $\text{dis}(\cdot)$ is a distance metric.

¹www.mathworks.com/help/matlab/ref/sylvester.html

Table 2: Statistics of the tested dataset in ZSL experiments.

| Dataset | aP&aY | AwA | CUB | SUN |
|------------------|--------|--------|--------|--------|
| # Seen classes | 20 | 40 | 150 | 707 |
| # Unseen classes | 12 | 10 | 50 | 10 |
| # Samples | 15,339 | 30,475 | 11,788 | 14,340 |
| # Attributes | 64 | 85 | 312 | 102 |

Cold-start Recommendation

Suppose that we use \mathbf{X} to denote the user behavior, e.g., purchase, browse and share, which in general is a user-item matrix. \mathbf{S} is the user attributes, e.g., user preferences, personal information and social network data. We can learn an encoder \mathbf{W} and a decoder \mathbf{W}^\top by Eq. (13). Then, for the new users, CSR aims to learn \mathbf{X}_{new} which indicates the potential user-item relationship, which can be achieved by:

$$\mathbf{X}_{new} = \mathbf{W}^\top \mathbf{S}_{new}. \quad (15)$$

At last, the recommendation will be formulated as a multi-label classification problem (Zhang and Zhou 2014). Specifically, we can use the logistic regression function to predict the recommendation probability of items, and recommend the top- k items.

Experiments

In this section, we verify the proposed method on both zero-shot recognition and cold-start recommendation tasks. From Eq. (13), we know that the main part of our method can be implemented by only one line of Matlab code. The complete codes will be released on publication.

Zero-shot Learning

Datasets. For zero-shot recognition, four most popular benchmarks are evaluated. For instance, **aPascal-aYahoo (aP&aY)** (Farhadi et al. 2009), **Animal with Attribute (AwA)** (Lampert, Nickisch, and Harmeling 2014), **SUN scene attribute dataset (SUN)** (Patterson and Hays 2012) and **Caltech-UCSD Birds-200-2011 (CUB)** (Wah et al. 2011). The statistics of the datasets are reported in Table 2.

Settings. For ZSL, we follow the experimental protocols reported in previous work (Ding, Shao, and Fu 2017; Kodirov, Xiang, and Gong 2017), deep convolutional neural networks (CNNs) features extracted by GoogLeNet (Szegedy et al. 2015), which is the 1024-dimensional activation of the final pooling layer, are used as input. The hyper-parameters λ and β are tuned by cross-validation using the training data.

Baselines. Five state-of-the-art work, e.g., DAP (Lampert, Nickisch, and Harmeling 2014), ESZSL (Romera-Paredes and Torr 2015), SSE (Zhang and Saligrama 2015), JLSE (Zhang and Saligrama 2016) and LESD (Ding, Shao, and Fu 2017), are selected as competitors.

Results and Discussions. The experimental results of ZSL are reported in Table 3. From the results, we can see that our approach, LLAE, performs much better than the compared methods. Since the compared methods cover a wide range of models which deploy different techniques for zero-shot learning, and the state-of-the-art method LESD (Ding,

Table 3: Accuracy (%) of zero-shot recognition. The best results are marked as bold numbers.

| Method | aP&aY | AwA | CUB | SUN | Avg. |
|--------|--------------|--------------|--------------|--------------|--------------|
| DAP | 38.23 | 60.51 | 39.14 | 71.92 | 52.45 |
| ESZSL | 24.37 | 75.31 | 48.75 | 82.12 | 57.64 |
| SSE | 46.22 | 76.35 | 30.49 | 82.51 | 58.89 |
| JLSE | 50.46 | 80.51 | 42.83 | 83.86 | 64.42 |
| LESD | 58.83 | 76.62 | 56.25 | 88.36 | 70.02 |
| Ours | 56.16 | 85.24 | 61.93 | 92.07 | 73.85 |

Shao, and Fu 2017) is reported recently, the performance of our method is quite favorable and significant.

Different from conventional ZSL methods, our model is bilateral. We consider not only the unilateral projection from the feature space to the attribute space but also the reconstruction from the attribute space to the feature space. Although our initial purpose of reconstruction is tailored for cold-start recommendation (one cannot generate the recommendation list without reconstruction from the user information in CSR), it benefits zero-shot recognition as well. The results verify our proposition of this work—ZSL and CSR share the similar problem framework, and the interdisciplinary study on them can benefit both of the communities.

Cold-start Recommendation

Datasets. For cold-start recommendation, we mainly use social data as side information. The following four datasets, which consist of image, video, blog and music recommendation, are used for evaluation.

- **Flickr** (Tang, Wang, and Liu 2012) is a dataset collected from *flickr.com*², which is a popular personal photos managing and sharing website. Users in flickr can tag photos and subscribe photos in terms of tags with which he is interested. For instance, a user can subscribe photos with tag “baseball”. The evaluated dataset consists of 80,513 users, 195 interest groups as the items, and a social network with 5,899,882 links.
- **BlogCatalog** (Tang, Wang, and Liu 2012) is a dataset collected from *blogcatalog.com*³, which is a popular blog collaboration system. Any article published by a blogger in blogcatalog can be cataloged into some groups according to the topics, e.g., “sports”, “business” and “technology”. The tested dataset consists of 10,312 users, 39 topics as items, and a social network with 333,983 links.
- **YouTube** (Tang, Wang, and Liu 2012) is a dataset collected from *youtube.com*⁴, which is a popular video watching and sharing website. Users in YouTube can also subscribe interested topics. The evaluated dataset consists of 1,138,499 users, 47 categories as items, and a social network with 2,990,443 links.
- **Hetrec11-LastFM** (Cantador, Brusilovsky, and Kuflik 2011) is a dataset collected from *last.fm*⁵, which is an on-

²<http://www.flickr.com>

³<http://www.blogcatalog.com>

⁴<http://www.youtube.com>

⁵<http://www.last.fm>

Table 4: CSR results of mAP@100 (%) on different datasets.

| Method | Flickr | BlogCatalog | YouTube | LastFM |
|---------|--------------|--------------|--------------|--------------|
| CBF-KNN | 28.05 | 32.71 | 34.21 | 17.12 |
| Cos-Cos | 31.42 | 41.06 | 46.67 | 12.26 |
| BPR-Map | 21.59 | 28.22 | 30.35 | 8.85 |
| CMF | 21.41 | 27.76 | 28.16 | 8.13 |
| LoCo | 33.57 | 45.35 | 48.79 | 18.09 |
| Ours | 39.25 | 50.13 | 52.45 | 23.07 |

line music system. Hetrec11-LastFM contains social networking, tagging, and music artist listening information. The tested dataset consists of 1,892 users, 17,632 artists as items, and 186,479 tag assignments.

Settings. For the evaluated datasets, we split each of them into two subsets, one includes 10% of the users as new users (test dataset) for cold-start, and the remainder of 90% users are collected as training data to learn the encoder and decoder. We deploy cross-validation with grid-search to tune all hyper-parameters on training data. Specifically, we select 80% users for training and 10% for validation. The new users are randomly selected, so we build 10 training-test folds and report the average results.

Following previous work (Sedhain et al. 2017), we deploy the widely used precision@k, recall@k and mean average precision (mAP@100) as the measurements. All the hyper-parameters in the objective are tuned by cross-validation. The following five previous work powered by different techniques are evaluated as baselines: CBF-KNN (Gantner et al. 2010), Cos-Cos (Sedhain et al. 2014), BPR-Map (Gantner et al. 2010), CMF (Krohn-Grimberghe et al. 2012) and LoCo (Sedhain et al. 2017).

Results and Discussions. The experimental results on different datasets are reported in Table 5 - Table 8. Table 4 shows the mAP@100 on the four datasets. Regarding to the experimental results, we have the following discussions:

1) It has been verified by ZSL that our bilateral formulation, i.e., projection from feature space to attribute space and reconstruction from attribute space to feature space, is very effective. In fact, the bilateral formulation is also the main reason that our model performs better in CSR tasks. As we stated in the section of related work, existing CSR methods generally learn a projection from the user behavior space to the user preference space. The learned projection is unilateral. In other words, the formulations of existing work focus on the mapping from the behavior to the preference, but it did not take reconstruction into consideration. If we take a further look at the CSR problem, we can find that the projection from user behavior to the user preference and the reconstruction from user preference to user behavior is equally important. The projection guarantees that warm users and cold users can be compared in the preference space. The reconstruction assures that the user behavior (recommendation list) can be generated from user preference.

2) Most of the baselines are based on matrix factorization (MF). The MF approaches, however, is highly sensitive and easily disturbed by noises when the observations are sparse. Compared with the recommendation scenarios where there




Figure 3: Parameters sensitivity (a-b) and convergence curve (c) of the proposed method on AwA dataset.




Figure 4: The effect of the low-rank and reconstruction constraint. Results of the CSR evaluations are the Precision@5 on different datasets. CUB dataset is used for ZSL.

are intensive feedbacks provided by users and only limited items, social network observations are extremely sparse. LoCo and our LLAE deploy low-rank representation to reveal the true data structure from corrupted inputs. Thus, they outperform the baselines in most evaluations.

3) Although LoCo deploys low-rank representation, it learns a unilateral projection. In social data facilitated CSR problems, there are a large number of users and only limited attributes. Thus, a user will be linked with several attributes, and a specific attribute will be attached with a lot of users. If we learn only one projection, trivial connections will be involved. For example, if we learn only one projection from user behavior to user attributes, a behavior can be linked with almost all of the attributes with a low weight since most of the elements in the projection matrix are not zero. Such a projection will weaken the true attributes, introduce over-fitting and relax the generalization ability. The reconstruction constraint in LLAE can automatically handle this because a bad projection will cause a bad reconstruction.

Model Analysis

Parameters Sensitivity. For different dataset, the hyper parameters vary from dataset to dataset. For instance, the optimal value of λ on some datasets is around 10^5 , while on others is less than 10. They, therefore, need to be selected by cross-validation. However, when λ is large, we also need to increase the value of β so that the effect of low-rank constraint would not be neglected (it is easy to be understood by referring to Eq. (13)). Since it is hard to illustrate the effects of the parameters on different datasets (the optimal values can differ by orders of magnitude), the effect of λ and β on AwA, as an example, are reported in Fig. 3(a) and Fig. 3(b).

Complexity and Convergence. It is worth noting that LLAE is a linear algorithm, the main part of LLAE can be implemented by only one line of Matlab code, and it runs faster than most of previous work. For instance, it only costs about 1/1000 training time of SSE on AwA. Since we update \mathbf{W} and \mathbf{V} in an iterative fashion, we show the convergence

Table 5: Cold-start recommendation results (%) on Flickr dataset.

| @k | BPR-Map | | CMF | | CBF-KNN | | Cos-Cos | | LoCo | | Ours | |
|----|-----------|--------|-----------|--------|-----------|--------|-----------|--------|-----------|--------|--------------|--------------|
| | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall |
| 1 | 16.99 | 13.52 | 15.91 | 12.34 | 20.62 | 15.79 | 22.33 | 17.58 | 28.64 | 22.52 | 33.61 | 26.34 |
| 5 | 7.29 | 28.12 | 7.55 | 27.24 | 10.10 | 36.69 | 10.34 | 37.56 | 11.99 | 43.47 | 15.32 | 48.27 |
| 10 | 4.85 | 37.18 | 5.12 | 36.57 | 6.58 | 47.13 | 6.71 | 47.84 | 7.25 | 51.62 | 11.03 | 55.74 |
| 20 | 3.28 | 50.21 | 3.28 | 46.67 | 4.01 | 57.17 | 4.12 | 58.24 | 4.12 | 58.31 | 7.67 | 62.45 |

Table 6: Cold-start recommendation results (%) on BlogCatalog dataset.

| @k | BPR-Map | | CMF | | CBF-KNN | | Cos-Cos | | LoCo | | Ours | |
|----|-----------|--------|-----------|--------|-----------|--------|-----------|--------|-----------|--------|--------------|--------------|
| | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall |
| 1 | 16.21 | 11.67 | 18.59 | 15.05 | 20.14 | 16.05 | 31.76 | 25.64 | 37.65 | 30.35 | 43.55 | 35.82 |
| 5 | 10.13 | 38.89 | 8.62 | 33.32 | 11.38 | 43.14 | 13.26 | 49.78 | 14.26 | 53.24 | 19.26 | 57.72 |
| 10 | 7.87 | 57.63 | 6.56 | 49.15 | 8.41 | 60.47 | 8.86 | 64.75 | 8.90 | 65.08 | 12.45 | 69.24 |
| 20 | 5.62 | 81.07 | 4.78 | 69.65 | 5.88 | 83.94 | 5.89 | 84.27 | 5.62 | 80.23 | 7.36 | 83.19 |

Table 7: Cold-start recommendation results (%) on YouTube dataset.

| @k | BPR-Map | | CMF | | CBF-KNN | | Cos-Cos | | LoCo | | Ours | |
|----|-----------|--------|-----------|--------|-----------|--------|-----------|--------|-----------|--------|--------------|--------------|
| | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall |
| 1 | 18.76 | 19.15 | 21.18 | 19.15 | 23.58 | 18.23 | 33.86 | 29.15 | 40.22 | 34.35 | 46.37 | 39.25 |
| 5 | 14.16 | 40.29 | 10.23 | 37.32 | 13.58 | 46.67 | 16.57 | 50.21 | 18.55 | 60.14 | 22.18 | 65.23 |
| 10 | 9.29 | 59.58 | 8.43 | 53.15 | 10.25 | 63.77 | 10.86 | 67.56 | 12.23 | 69.34 | 15.39 | 73.25 |
| 20 | 6.25 | 83.51 | 5.91 | 73.25 | 6.67 | 84.61 | 6.84 | 86.59 | 8.77 | 88.74 | 11.32 | 92.58 |

Table 8: Cold-start recommendation results (%) on Hetrec11-LastFM dataset.

| @k | BPR-Map | | CMF | | CBF-KNN | | Cos-Cos | | LoCo | | Ours | |
|----|-----------|--------|-----------|--------|-----------|--------|-----------|--------|-----------|--------|--------------|--------------|
| | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall | Precision | Recall |
| 1 | 32.53 | 0.81 | 37.54 | 0.94 | 53.39 | 1.35 | 36.54 | 0.92 | 58.11 | 1.42 | 62.95 | 4.33 |
| 5 | 26.83 | 3.07 | 30.71 | 3.54 | 45.08 | 5.14 | 31.97 | 3.87 | 48.09 | 5.51 | 55.16 | 10.27 |
| 10 | 23.99 | 5.31 | 25.64 | 5.95 | 39.65 | 8.91 | 28.07 | 6.56 | 41.87 | 9.58 | 46.67 | 13.19 |
| 20 | 20.53 | 9.05 | 20.97 | 9.67 | 32.32 | 14.39 | 22.95 | 10.56 | 33.98 | 15.42 | 39.86 | 18.45 |

curve of LLAE on AwA dataset in Fig. 3(c). It can be seen that our model converges very fast.

Low-rank Constraint. From the parameter curve of β , we can see that the low-rank constraint is effective for the performance of LLAE. For ZSL, it helps to find out shared semantics across different categories. For CSR, it filters out spurious connections and handles the extremely sparse observations. If we investigate the CSR as a specific transfer learning problem (Ding, Shao, and Fu 2018; Li et al. 2018c), the low-rank constraint can also mitigate the domain shift between the user behavior space and the user attribute space. We show the effects of the low-rank constraints in Fig. 4(a).

Reconstruction Constraint. We have mentioned in several places that the reconstruction constraint plays an important role in our formulation. To verify the claim, Fig. 4(b) shows the effects of the reconstruction part on several evaluations. It is obvious that the reconstruction constraint contributes a lot for the performance.

Conclusion

This paper, for the first time, investigates CSR as a ZSL problem. Although CSR and ZSL were independently studied by two communities in general, we reveal that the two

tasks are two extensions of the same intension. In this paper, a tailor-made ZSL model is proposed to handle CSR. Specifically, we present a low-rank linear autoencoder, which deploys a low-rank encoder to map user behavior space to user attribute space and a symmetric decoder to reconstruct user behavior from the user attributes. Extensive experiments on eight datasets, including both CSR and ZSL, verify not only that the CSR problem can be addressed by ZSL model, but the consideration of CSR, e.g., the reconstruction constraint, can benefit ZSL as well. It is a win-win formulation. At last, by linking CSR and ZSL, we wish that this work will benefit both of the communities and elicit more contributions. In our future work, we are going to investigate training deep autoencoders for both ZSL and CSR.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 61806039, 61832001, 61802236, 61572108 and 61632007, in part by the ARC under Grant FT130101530, in part by the National Postdoctoral Program for Innovative Talents under Grant BX201700045, and in part by the China Postdoctoral Science Foundation under Grant 2017M623006.

References

- Bobadilla, J.; Ortega, F.; Hernando, A.; and Gutiérrez, A. 2013. Recommender systems survey. *Knowledge-based systems* 46:109–132.
- Boureau, Y.-L.; Chopra, S.; LeCun, Y.; et al. 2007. A unified energy-based framework for unsupervised learning. In *Artificial Intelligence and Statistics*, 371–379.
- Cantador, I.; Brusilovsky, P. L.; and Kuflik, T. 2011. Second workshop on information heterogeneity and fusion in recommender systems (hetrec2011).
- Ding, Z.; Shao, M.; and Fu, Y. 2017. Low-rank embedded ensemble semantic dictionary for zero-shot learning. In *CVPR*. IEEE.
- Ding, Z.; Shao, M.; and Fu, Y. 2018. Incomplete multisource transfer learning. *IEEE TNNLS* 29(2):310–323.
- Ekstrand, M. D.; Riedl, J. T.; Konstan, J. A.; et al. 2011. Collaborative filtering recommender systems. *Foundations and Trends® in Human-Computer Interaction* 4(2):81–173.
- Farhadi, A.; Endres, I.; Hoiem, D.; and Forsyth, D. 2009. Describing objects by their attributes. In *CVPR*, 1778–1785. IEEE.
- Fernández-Tobías, I.; Cantador, I.; Kaminskis, M.; and Ricci, F. 2012. Cross-domain recommender systems: A survey of the state of the art. In *Spanish Conference on Information Retrieval*, 24.
- Fernández-Tobías, I.; Braunhofer, M.; Elahi, M.; Ricci, F.; and Cantador, I. 2016. Alleviating the new user problem in collaborative filtering by exploiting personality information. *User Modeling and User-Adapted Interaction* 26(2-3):221–255.
- Gantner, Z.; Drumond, L.; Freudenthaler, C.; Rendle, S.; and Schmidt-Thieme, L. 2010. Learning attribute-to-feature mappings for cold-start recommendations. In *ICDM*, 176–185. IEEE.
- Kodirov, E.; Xiang, T.; Fu, Z.; and Gong, S. 2015. Unsupervised domain adaptation for zero-shot learning. In *ICCV*, 2452–2460.
- Kodirov, E.; Xiang, T.; and Gong, S. 2017. Semantic autoencoder for zero-shot learning. *arXiv preprint arXiv:1704.08345*.
- Krohn-Grimberghe, A.; Drumond, L.; Freudenthaler, C.; and Schmidt-Thieme, L. 2012. Multi-relational matrix factorization using bayesian personalized ranking for social network data. In *ACM WSDM*, 173–182. ACM.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2014. Attribute-based classification for zero-shot visual object categorization. *IEEE TPAMI* 36(3):453–465.
- Li, J.; Wu, Y.; Zhao, J.; and Lu, K. 2016. Low-rank discriminant embedding for multiview learning. *IEEE TCYB*.
- Li, J.; Lu, K.; Huang, Z.; and Shen, H. T. 2017. Two birds one stone: On both cold-start and long-tail recommendation. In *ACM MM*, 898–906. ACM.
- Li, J.; Lu, K.; Huang, Z.; Zhu, L.; and Shen, H. T. 2018a. Heterogeneous domain adaptation through progressive alignment. *IEEE TNNLS*.
- Li, J.; Lu, K.; Huang, Z.; Zhu, L.; and Shen, H. T. 2018b. Transfer independently together: A generalized framework for domain adaptation. *IEEE TCYB*.
- Li, J.; Zhu, L.; Huang, Z.; Lu, K.; and Zhao, J. 2018c. I read, i saw, i tell: Texts assisted fine-grained visual classification. In *ACM MM*, 663–671. ACM.
- Lin, J.; Sugiyama, K.; Kan, M.-Y.; and Chua, T.-S. 2013. Addressing cold-start in app recommendation: latent user models constructed from twitter followers. In *ACM SIGIR*, 283–292. ACM.
- Mohamed, A.-r.; Dahl, G. E.; and Hinton, G. 2012. Acoustic modeling using deep belief networks. *IEEE TASLP* 20(1):14–22.
- Noel, J.; Sanner, S.; Tran, K.-N.; Christen, P.; Xie, L.; Bonilla, E. V.; Abbasnejad, E.; and Della Penna, N. 2012. New objective functions for social collaborative filtering. In *WWW*, 859–868. ACM.
- Patterson, G., and Hays, J. 2012. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, 2751–2758. IEEE.
- Rohani, V. A.; Kasirun, Z. M.; Kumar, S.; and Shamshirband, S. 2014. An effective recommender algorithm for cold-start problem in academic social networks. *Mathematical Problems in Engineering* 2014.
- Romera-Paredes, B., and Torr, P. 2015. An embarrassingly simple approach to zero-shot learning. In *ICML*, 2152–2161.
- Sedhain, S.; Sanner, S.; Brazianus, D.; Xie, L.; and Christensen, J. 2014. Social collaborative filtering for cold-start recommendations. In *ACM RecSys*, 345–348. ACM.
- Sedhain, S.; Menon, A. K.; Sanner, S.; Xie, L.; and Brazianus, D. 2017. Low-rank linear cold-start recommendation from social data. In *AAAI*, 1502–1508.
- Smith, B., and Linden, G. 2017. Two decades of recommender systems at amazon.com. *IEEE Internet Computing* 21(3):12–18.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *CVPR*, 1–9.
- Tang, L.; Wang, X.; and Liu, H. 2012. Scalable learning of collective behavior. *TKDE* 24(6):1080–1091.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P. A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*, 1096–1103.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Yang, Y.; Luo, Y.; Chen, W.; Shen, F.; Shao, J.; and Shen, H. T. 2016. Zero-shot hashing via transferring supervised knowledge. In *ACM MM*, 1286–1295. ACM.
- Zhang, Z., and Saligrama, V. 2015. Zero-shot learning via semantic similarity embedding. In *ICCV*, 4166–4174.
- Zhang, Z., and Saligrama, V. 2016. Zero-shot learning via joint latent similarity embedding. In *CVPR*, 6034–6042.
- Zhang, M.-L., and Zhou, Z.-H. 2014. A review on multi-label learning algorithms. *IEEE TKDE* 26(8):1819–1837.