

Deep Adversarial Social Recommendation

Wenqi Fan¹, Tyler Derr², Yao Ma², Jianping Wang¹, Jiliang Tang² and Qing Li³

¹Department of Computer Science, City University of Hong Kong

²Data Science and Engineering Lab, Michigan State University

³Department of Computing, The Hong Kong Polytechnic University

wenqifan03@gmail.com, {derrtyl, mayao4}@msu.edu, jianwang@cityu.edu.hk, tangjili@msu.edu, csqli@comp.polyu.edu.hk

Abstract

Recent years have witnessed rapid developments on social recommendation techniques for improving the performance of recommender systems due to the growing influence of social networks to our daily life. The majority of existing social recommendation methods unify user representation for the user-item interactions (item domain) and user-user connections (social domain). However, it may restrain user representation learning in each respective domain, since users behave and interact differently in two domains, which makes their representations to be heterogeneous. In addition, most of traditional recommender systems can not efficiently optimize these objectives, since they utilize negative sampling technique which is unable to provide enough informative guidance towards the training during the optimization process. In this paper, to address the aforementioned challenges, we propose a novel **Deep Adversarial SO**cial recommendation **DASO**. It adopts a bidirectional mapping method to transfer users' information between social domain and item domain using adversarial learning. Comprehensive experiments on two real-world datasets show the effectiveness of the proposed method.

1 Introduction

In recent years, we have seen an increasing amount of attention on social recommendation, which harnesses social relations to boost the performance of recommender systems [Tang *et al.*, 2016b; Fan *et al.*, 2019; Wang *et al.*, 2016]. Social recommendation is based on the intuitive ideas that people in the same social group are likely to have similar preferences, and that users will gather information from their experienced friends (e.g., classmates, relatives, and colleagues) when making decisions. Therefore, utilizing users' social relations has been proven to greatly enhance the performance of many recommender systems [Ma *et al.*, 2008; Fan *et al.*, 2019; Tang *et al.*, 2013b; 2016a].

In Figure 1, we observe that in social recommendation we have both the item and social domains, which represent the user-item interactions and user-user connections, respec-




Figure 1: An illustration of one user in two domains (Item Domain and Social Domain) for social recommendations.

tively. Currently, the most effective way to incorporate the social information for improving recommendations is when learning user representations, which is commonly achieved in ways such as, using trust propagation [Jamali and Ester, 2010], incorporating a user's social neighborhood information [Fan *et al.*, 2018], or sharing a common user representation for the user-item interactions and social relations with a co-factorization method [Ma *et al.*, 2008]. However, as shown in Figure 1, although users bridge the gap between these two domains, their representations should be heterogeneous. This is because users behave and interact differently in the two domains. Thus, using a unified user representation may restrain user representation learning in each respective domain and results in an inflexible/limited transferring of knowledge from the social relations to the item domain. Therefore, one challenge is to learn separated user representations in two domains while transferring the information from the social domain to the item domain for recommendation.

In this paper, we adopt a nonlinear mapping operation to transfer user's information from the social domain to the item domain, while learning separated user representations in the two domains. Nevertheless, learning the representations is challenging due to the inherent data sparsity problem in both domains. Thus, to alleviate this problem, we propose to use a bidirectional mapping between the two domains, such that we can cycle information between them to progressively enhance the user's representations in both domains. However, for optimizing the user representations and item representations, most existing methods utilize the negative sampling technique, which is quite ineffective [Wang *et al.*, 2018b]. This is due to the fact that during the beginning of the training process, most of the negative user-item sam-

ples are still within the margin to the real user-item samples, but later during the optimization process, negative sampling is unable to provide “difficult” and informative samples to further improve the user representations and item representations [Wang *et al.*, 2018b; Cai and Wang, 2018]. Thus, it is desired to have samples dynamically generated throughout the training process to better guide the learning of the user representations and item representations.

Recently, Generative Adversarial Networks (GANs) [Goodfellow *et al.*, 2014], which consists of two models to process adversarial learning, have shown great success across various domains due to their ability to learn an underlying data distribution and generate synthetic samples [Mao *et al.*, 2017; 2018; Brock *et al.*, 2019; Liu *et al.*, 2018; Wang *et al.*, 2017; 2018a; Derr *et al.*, 2019]. This is performed through the use of a generator and a discriminator. The generator tries to generate realistic fake data samples to fool the discriminator, which distinguishes whether a given data sample is produced by the generator or comes from the real data distribution. A minimax game is played between the generator and discriminator, where this adversarial learning can train these two models simultaneously for mutual promotion. In [Wang *et al.*, 2018b] adversarial learning had been used to address the limitation of typical negative sampling. Thus, we propose to harness adversarial learning in social recommendation to generate “difficult” negative samples to guide our framework in learning better user and item representations while further utilizing it to optimize our entire framework. Our major contributions can be summarized as follows:

- We introduce a principled way to transfer users’ information from social domain to item domain using a bidirectional mapping method where we cycle information between the two domains to progressively enhance the user representations;
- We propose a Deep Adversarial SOcial recommender system **DASO**, which can harness the power of adversarial learning to dynamically generate “difficult” negative samples, learn the bidirectional mappings between the two domains, and ultimately optimize better user and item representations; and
- We conduct comprehensive experiments on two real-world datasets to show the effectiveness of the proposed model.

2 The Proposed Framework

Let $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$ and $\mathcal{V} = \{v_1, v_2, \dots, v_M\}$ denote the sets of users and items respectively, where N (M) is the number of users (items). We define user-item interactions matrix $\mathbf{R} \in \mathbb{R}^{N \times M}$ from user’s implicit feedback, where the i, j -th element $r_{i,j}$ is 1 if there is an interaction (e.g., clicked/bought) between user u_i and item v_j , and 0 otherwise. However, $r_{i,j} = 1$ does not mean user u_i actually likes item v_j . Similarly, $r_{i,j} = 0$ does not mean u_i does not like item v_j , since it can be that the user u_i is not aware of the item v_j . The social network between users can be described by a matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$, where $s_{i,j} = 1$ if there is a social relation between user u_i and user u_j , and 0 otherwise. Given interactions matrix \mathbf{R} and social network \mathbf{S} , we aim to predict

the unobserved entries (i.e., those where $r_{i,j} = 0$) in \mathbf{R} .

2.1 An Overview of the Proposed Framework

The architecture of the proposed model is shown in Figure 2. The information is from two domains, which are the item domain I and the social domain S . The model consists of three components: cyclic user modeling, item domain adversarial learning, and social domain adversarial learning. The cyclic user modeling is to model user representations on two domains. The item domain adversarial learning is to adopt the adversarial learning for dynamically generating “difficult” and informative negative samples to guide the learning of user and item representations. The generator is utilized to ‘sample’ (recommend) items for each user and output user-item pairs as fake samples; the other is the discriminator, which distinguishes the user-item pair samples sampled from the real user-item interactions from the generated user-item pair samples. The social domain adversarial learning also similarly consists of a generator and a discriminator.

There are four types of representations in the two domains. In the item domain I , we have two types of representations including item domain representations of the generator ($\mathbf{p}_i^I \in \mathbb{R}^d$ for user u_i and $\mathbf{q}_j^I \in \mathbb{R}^d$ for item v_j), and the item domain representations of the discriminator ($\mathbf{x}_i^I \in \mathbb{R}^d$ for user u_i and $\mathbf{y}_j^I \in \mathbb{R}^d$ for item v_j). Social domains S also contains two types of representations including the social domain representations of the generator ($\mathbf{p}_i^S \in \mathbb{R}^d$ for user u_i), and the social domain representations of the discriminator ($\mathbf{x}_i^S \in \mathbb{R}^d$ for user u_i).

2.2 Cyclic User Modeling

Cyclic user modeling aims to learn a relation between the user representations in the item domain I and the social domain S . As shown in the top part of Figure 2, we first adopt a nonlinear mapping operation, denoted as $h^{S \rightarrow I}$, to transfer user’s information from the social domain to the item domain, while learning separated user representations in the two domains. Then, a bidirectional mapping between these two domains (achieved by including another nonlinear mapping $h^{I \rightarrow S}$) is utilized to help cycle the information between them to progressively enhance the user representations in both domains.

Transferring Social Information to Item Domain

In social networks, a person’s preferences can be influenced by their social interactions, suggested by sociologists [Fan *et al.*, 2019; 2018; Wasserman and Faust, 1994]. Therefore, a user’s social relations from the social network should be incorporated into their user representation in the item domain.

We propose to adopt nonlinear mapping operation to transfer user’s information from the social domain to the item domain. More specifically, the user representation on social domain \mathbf{p}_i^S is transferred to the item domain via a Multi-Layer Perceptron (MLP) denoted as $h^{S \rightarrow I}$. The transferred user representation from social domain is denoted as \mathbf{p}_i^{SI} . More formally, the nonlinear mapping is as follows: $\mathbf{p}_i^{SI} = h^{S \rightarrow I}(\mathbf{p}_i^S) = \mathbf{W}_L \cdot (\dots a(\mathbf{W}_2 \cdot a(\mathbf{W}_1 \cdot \mathbf{p}_i^S + \mathbf{b}_1) + \mathbf{b}_2) \dots) + \mathbf{b}_L$, where the \mathbf{W} s, \mathbf{b} s are the weights and biases for the layers of the neural network having L layers, and a is a nonlinear activation function.

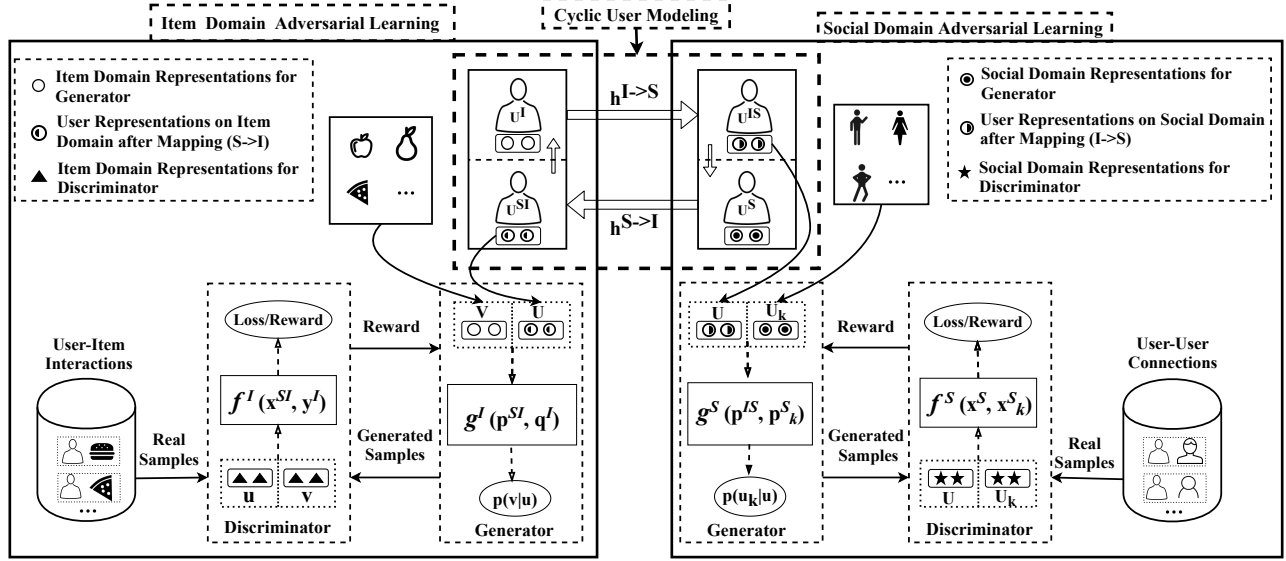


Figure 2: The overall architecture of the proposed model DASO.

Bidirectional Mapping with Cycle Reconstruction

As user-item interactions and user-user connections are often very sparse, learning separated user representations is challenging. Therefore, to partially alleviate this issue, we propose to utilize a bidirectional mapping between the two domains, such that we can cycle information between them to progressively enhance the user representations in both domains. To achieve this, another nonlinear mapping operation, denoted as $h^{I \rightarrow S}$, is adopted to transfer information from the item domain to the social domain: $\mathbf{p}_i^{IS} = h^{I \rightarrow S}(\mathbf{p}_i^I)$, which has the same network structure as the $h^{S \rightarrow I}$.

This Bidirectional Mapping allows knowledge to be transferred between item and social domains. To learn these mappings, we further introduce cycle reconstruction. Its intuition is that transferred knowledge in the target domain should be reconstructed to the original knowledge in the source domain. Next we will elaborate cycle reconstruction.

For user u_i 's item domain representation \mathbf{p}_i^I , the user representation with cycle reconstruction should be able to map \mathbf{p}_i^I back to the original domain, as follows, $\mathbf{p}_i^I \rightarrow h^{I \rightarrow S}(\mathbf{p}_i^I) \rightarrow h^{S \rightarrow I}(h^{I \rightarrow S}(\mathbf{p}_i^I)) \approx \mathbf{p}_i^I$. Likewise, for user u_i 's social domain representation \mathbf{p}_i^S , the user representation with cycle reconstruction can also bring \mathbf{p}_i^S back to the original domain: $\mathbf{p}_i^S \rightarrow h^{S \rightarrow I}(\mathbf{p}_i^S) \rightarrow h^{I \rightarrow S}(h^{S \rightarrow I}(\mathbf{p}_i^S)) \approx \mathbf{p}_i^S$.

We can formulate this procedure using a cycle reconstruction loss, which needs to be minimized, as follows,

$$\mathcal{L}_{cyc}(h^{S \rightarrow I}, h^{I \rightarrow S}) = \sum_{i=1}^N (\|h^{S \rightarrow I}(h^{I \rightarrow S}(\mathbf{p}_i^I)) - \mathbf{p}_i^I\|_2 + \|h^{I \rightarrow S}(h^{S \rightarrow I}(\mathbf{p}_i^S)) - \mathbf{p}_i^S\|_2)$$

2.3 Item Domain Adversarial Learning

To address the limitation of negative sampling for recommendation on the ranking task, we propose to harness adversarial learning to generate “difficult” and informative samples to

guide the framework in learning better user and item representations in the item domain. As shown in the bottom left part of Figure 2, the adversarial learning on item domain consists of two components:

Discriminator $D^I(u_i, v; \phi_D^I)$, parameterized by ϕ_D^I , aims to distinguish the real user-item pairs (u_i, v) and the user-item pairs generated by the generator.

Generator $G^I(v|u_i; \theta_G^I)$, parameterized by θ_G^I , tries to fit the underlying real conditional distribution $p_{real}^I(v|u_i)$ as much as possible, and generates (or, to be more precise, selects) the most relevant items to a given user u_i .

Formally, D^I and G^I are playing the following two-player minimax game with value function $\mathcal{L}_{adv}^I(G^I, D^I)$,

$$\begin{aligned} \min_{\theta_G^I} \max_{\phi_D^I} \mathcal{L}_{adv}^I(G^I, D^I) \\ = \sum_{i=1}^N (\mathbb{E}_{v \sim p_{real}^I(\cdot|u_i)} [\log D^I(u_i, v; \phi_D^I)] \\ + \mathbb{E}_{v \sim G^I(\cdot|u_i; \theta_G^I)} [\log(1 - D^I(u_i, v; \phi_D^I))]) \end{aligned} \quad (1)$$

Item Domain Discriminator Model

Discriminator D^I aims to distinguish real user-item pairs (i.e., real samples) and the generated “fake” samples. The discriminator D^I estimates the probability of item v_j being relevant (bought or clicked) to a given user u_i using the sigmoid function and a score function $f_{\phi_D^I}^I$ as follows:

$$D^I(u_i, v_j; \phi_D^I) = \frac{1}{1 + \exp(-f_{\phi_D^I}^I(\mathbf{x}_i^I, \mathbf{y}_j^I))} \quad (2)$$

Given real samples and generated fake samples, the objective for the discriminator D^I is to maximize the log-likelihood of assigning the correct labels to both real and generated samples. The discriminator can be optimized by minimizing the objective in eq. (1) with the generator fixed using stochastic gradient methods.

Item Domain Generator Model

On the other hand, the purpose of the generator G^I is to approximate the underlying real conditional distribution $p_{real}^I(v|u_i)$, and generate the most relevant items for any given user u_i .

We define the generator using the softmax function over all the items according to the transferred user representation \mathbf{p}_i^{SI} from social domain to item domain:

$$G^I(v_j|u_i; \theta_G^I) = \frac{\exp(g_{\theta_G^I}^I(\mathbf{p}_i^{SI}, \mathbf{q}_j^I))}{\sum_{v_j \in \mathcal{V}} \exp(g_{\theta_G^I}^I(\mathbf{p}_i^{SI}, \mathbf{q}_j^I))} \quad (3)$$

where $g_{\theta_G^I}^I$ is a score function reflecting the chance of v_j being clicked/purchased by u_i . Given a user u_i , an item v_j can be sampled from the distribution $G^I(v_j|u_i; \theta_G^I)$.

We note that the process of generating a relevant item for a given user is discrete. Thus, we cannot optimize the generator G^I via stochastic gradient descent methods [Wang *et al.*, 2017]. Following [Sutton *et al.*, 2000; Schulman *et al.*, 2015], we adopt the policy gradient method usually adopted in reinforcement learning to optimize the generator.

To learn the parameters for the generator, we need to perform the following minimization problem:

$$\begin{aligned} & \min_{\theta_G^I} \sum_{i=1}^N \left(\mathbb{E}_{v \sim G^I(\cdot|u_i; \theta_G^I)} [\log(1 - D^I(u_i, v; \phi_D^I))] \right) \quad (4) \\ & = \max_{\theta_G^I} \sum_{i=1}^N \left(\mathbb{E}_{v \sim G^I(\cdot|u_i; \theta_G^I)} [\log(1 + \exp(f_{\phi_D^I}^I(\mathbf{x}_i^I, \mathbf{y}_j^I)))] \right) \end{aligned}$$

Now, this problem can be viewed in a reinforcement learning setting, where $K(\mathbf{x}_i^I, \mathbf{y}_j^I) = \log(1 + \exp(f_{\phi_D^I}^I(\mathbf{x}_i^I, \mathbf{y}_j^I)))$ is the reward given to the action ‘‘selecting v_i given a user u_i ’’ performed according to the policy probability $G^I(v|u_i)$. The policy gradient can be written as:

$$\begin{aligned} \nabla_{\theta_G^I} \mathcal{L}_{adv}^I(G^I, D^I) &= \sum_{i=1}^N \sum_{j=1}^M \nabla_{\theta_G^I} G^I(v_j|u_i) K(\mathbf{x}_i^I, \mathbf{y}_j^I) \\ &= \sum_{i=1}^N \sum_{j=1}^M G^I(v_j|u_i) \nabla_{\theta_G^I} \log G^I(v_j|u_i) K(\mathbf{x}_i^I, \mathbf{y}_j^I) \\ &= \sum_{i=1}^N \mathbb{E}_{v_j \sim G^I(\cdot|u_i)} \left[\nabla_{\theta_G^I} \log G^I(v_j|u_i) K(\mathbf{x}_i^I, \mathbf{y}_j^I) \right] \quad (5) \end{aligned}$$

Specially, the gradient $\nabla_{\theta_G^I} \mathcal{L}_{adv}^I(G^I, D^I)$ is an expected summation over the gradients $\nabla_{\theta_G^I} \log G^I(v_j|u_i)$ weighted by $\log(1 + \exp(f_{\phi_D^I}^I(\mathbf{x}_i^I, \mathbf{y}_j^I)))$.

The optimal parameters of G^I and D^I can be learned by alternately minimizing and maximizing the value function $\mathcal{L}_{adv}^I(G^I, D^I)$. In each iteration, discriminator D^I is trained with real samples from $p_{real}^I(\cdot|u_i)$ and generated samples from generator G^I ; the generator G^I is updated with policy gradient under the guidance of D^I .

Note that different from the way of optimizing user and item representations with the typical negative sampling on traditional recommender systems, the adversarial learning technique tries to generate ‘‘difficult’’ and high-quality negative samples to guide the learning of user and item representations.

2.4 Social Domain Adversarial Learning

In order to learn better user representations from the social perspective, another adversarial learning is harnessed in the social domain. Likewise, the adversarial learning in the social domain consists of two components, as shown in the bottom right part of Figure 2.

Discriminator $D^S(u_i, u; \phi_D^S)$, parameterized by ϕ_D^S , aims to distinguish the real connected user-user pairs (u_i, u) and the fake user-user pairs generated by the generator G^S .

Generator $G^S(u|u_i; \theta_G^S)$, parameterized by θ_G^S , tries to fit the underlying real conditional distribution $p_{real}^S(u|u_i)$ as much as possible, and generates (or, to be more precise, selects) the most relevant users to the given user u_i .

Formally, D^S and G^S are playing the following two-player minimax game with value function $\mathcal{L}_{adv}^S(G^S, D^S)$,

$$\begin{aligned} & \min_{\theta_G^S} \max_{\phi_D^S} \mathcal{L}_{adv}^S(G^S, D^S) \quad (6) \\ & = \sum_{i=1}^N \left(\mathbb{E}_{u \sim p_{real}^S(\cdot|u_i)} [\log D^S(u_i, u; \phi_D^S)] \right. \\ & \quad \left. + \mathbb{E}_{u \sim G^S(\cdot|u_i; \theta_G^S)} [\log(1 - D^S(u_i, u; \phi_D^S))] \right) \end{aligned}$$

Social Domain Discriminator

The discriminator D^S aims to distinguish the real user-user pairs and the generated ones. The discriminators D^S estimates the probability of user u_k being connected to user u_i with a sigmoid function and a score function $f_{\phi_D^S}^S$ as follows:

$$D^S(u_i, u_k; \phi_D^S) = \frac{1}{1 + \exp(-f_{\phi_D^S}^S(\mathbf{x}_i^S, \mathbf{x}_k^S))} \quad (7)$$

Social Domain Generator

The purpose of the generator, G^S , is to approximate the underlying real conditional distribution $p_{real}^S(u|u_i)$, and generate (or, to be more precise, select) the most relevant users for any given user u_i .

We model the distribution using a softmax function over all the other users with the transferred user representation \mathbf{p}_i^{IS} (from the item to social domain),

$$G^S(u_k|u_i; \theta_G^S) = \frac{\exp(g_{\theta_G^S}^S(\mathbf{p}_i^{IS}, \mathbf{p}_k^S))}{\sum_{u_k \neq u_i} \exp(g_{\theta_G^S}^S(\mathbf{p}_i^{IS}, \mathbf{p}_k^S))} \quad (8)$$

where $g_{\theta_G^S}^S$ is a score function reflecting the chance of u_k being related to u_i .

Likewise, policy gradient is utilized to optimize the generator G^S ,

$$\begin{aligned} \nabla_{\theta_G^S} \mathcal{L}_{adv}^S(G^S, D^S) &= \sum_{i=1}^N \mathbb{E}_{u_k \sim G^S(\cdot|u_i)} \\ & \left[\nabla_{\theta_G^S} \log G^S(u_k|u_i) \log(1 + \exp(f_{\phi_D^S}^S(\mathbf{x}_i^S, \mathbf{x}_k^S))) \right] \quad (9) \end{aligned}$$

Datasets	Ciao	Epinions
# of Users	7,317	14,575
# of Items	10,4975	155,527
# of Interactions	283,319	418,936
Density of Interactions	0.0368%	0.0184%
# of Social Relations	111,781	249,586
Density of Social Relations	0.2087%	0.1175%

Table 1: Statistics of the datasets.

where the details are omitted here, since it is defined similar to Eq.(5).

2.5 The Objective Function

With all model components, the objective function of the proposed framework is:

$$\begin{aligned} & \min_{G^I, G^S, h^{S \rightarrow I}, h^{I \rightarrow S}} \max_{D^I, D^S} \mathcal{L} \\ & = \mathcal{L}_{adv}^I(G^I, D^I) + \mathcal{L}_{adv}^S(G^S, D^S) + \lambda \mathcal{L}_{cyc}(h^{S \rightarrow I}, h^{I \rightarrow S}) \end{aligned} \quad (10)$$

where λ is to control the relative importance of cycle-reconstruction strategy and further influences the two mapping operation. $h^{S \rightarrow I}$ and $h^{I \rightarrow S}$ are implemented as MLP with three hidden layers. To optimize the objective, the RM-Sprop [Tieleman and Hinton, 2012] is adopted as the optimizer in our implementation. To train our model, at each training epoch, we iterate over the training set in mini-batch to train each model (e.g., G^I) while the parameters of other models (e.g., D^I, G^S, D^S) are fixed. When the training is finished, we take the representations learned by the generator G^I and G^S as our final representations of item and user for performing recommendation.

There are six representations in our model, including $\mathbf{p}_i^I, \mathbf{q}_j^I, \mathbf{x}_i^I, \mathbf{y}_j^I, \mathbf{p}_i^S, \mathbf{x}_i^S$. They are randomly initialized and jointly learned during the training stage.

Following the setting of IRGAN [Wang *et al.*, 2017], we adopt the inner product as the score function $f_{\phi_D^I}^I$ and $g_{\theta_G^I}^I$ in the item domain as follows: $f_{\phi_D^I}^I(\mathbf{x}_i^I, \mathbf{y}_j^I) = (\mathbf{x}_i^I)^T \mathbf{y}_j^I + a_j$, $g_{\theta_G^I}^I(\mathbf{p}_i^{SI}, \mathbf{q}_j^I) = (\mathbf{p}_i^{SI})^T \mathbf{q}_j^I + b_j$, where a_j and b_j are the bias term for item j . We define the score function $f_{\phi_D^S}$ and $g_{\theta_G^S}$ in the social domain in a similar way. Note that the above score functions can be also implemented using deep neural networks, but leave this investigation as one future work.

3 Experiments

3.1 Experimental Settings

We conduct our experiments on two representative datasets Ciao and Epinions¹ for the Top-K recommendation. As these two datasets provide users' explicit ratings on items, we convert them into 1 as the implicit feedback. This processing method is widely used in previous works on recommendation with implicit feedback [Rendle *et al.*, 2009]. We randomly split the user-item interactions of each dataset into training set (80%) to learn the parameters, validation set (10%) to tune

¹Both Ciao and Epinions datasets are available at: <http://www.cse.msu.edu/~tangjili/trust.html>

hyper-parameters, and testing set (10%) for the final performance comparison. We implemented our method with tensor-flow and tuned all the hyper-parameters with grid-search [Fan *et al.*, 2019]. The statistics of these two datasets are presented in Table 1. We use two popular performance metrics for Top-K recommendation [Wang *et al.*, 2017]: Precision@K and Normalized Discounted Cumulative Gain (NDCG@K). We set K as 3, 5, and 10. Higher values of the Precision@K and NDCG@K indicate better predictive performance.

To evaluate the performance, we compared our proposed model **DASO** with four groups of representative baselines, including traditional recommender system without social network information (**BPR** [Rendle *et al.*, 2009]), traditional social recommender systems (**SBPR** [Zhao *et al.*, 2014] and **SocialMF** [Jamali and Ester, 2010]), deep neural networks based social recommender systems (**DeepSoR** [Fan *et al.*, 2018] and **GraphRec** [Fan *et al.*, 2019]), and adversarial learning based recommender system (**IRGAN** [Wang *et al.*, 2017]). Some of the original baseline implementations (SocialMF, DeepSoR, and GraphRec) are for rating prediction on recommendations. Therefore we adjust their objectives to point-wise prediction with sigmoid cross entropy loss using negative sampling.

3.2 Performance Comparison of Recommender Systems

Table 2 presents the performance of all recommendation methods. We have the following findings:

- SBPR and SocialMF outperform BPR. SBPR and SocialMF utilize both user-item interactions and social relations; while BPR only uses the user-item interactions. These improvements show the effectiveness of incorporating social relations for recommender systems.
- In most cases, the two deep models, DeepSoR and GraphRec, obtain better performance than SBPR and SocialMF, which are modeled with shallow architectures. These improvements reflect the power of deep architectures on the task of recommendations.
- IRGAN achieves much better performance than BPR, while both of them utilize the user-item interactions only. IRGAN adopts the adversarial learning to optimize user and item representations; while BPR is a pair-wise ranking framework for Top-K traditional recommender systems. This suggests that adopting adversarial learning can provide more informative negative samples and thus improve the performance of the model.
- Our model DASO consistently outperforms all the baselines. Compared with DeepSoR and GraphRec, our model proposes advanced model components to model user representations in both item domain and social domain. In addition, our model harnesses the power of adversarial learning to generate more informative negative samples, which can help learn better user and item representations.

Parameter Analysis

Next, we investigate how the value of λ affects the performance of the proposed framework. The value of λ is to control the importance of cycle reconstruction. Figure 3 shows

Datasets	Metrics	Algorithms						
		BPR	IRGAN	SBPR	SocialMF	DeepSoR	GraphRec	DASO
Ciao	Precision@3	0.0154	0.0274	0.0211	0.0260	0.0310	0.0374	0.0462
	Precision@5	0.0137	0.0245	0.0204	0.0218	0.0240	0.0326	0.0451
	Precision@10	0.0102	0.0239	0.0178	0.0155	0.0201	0.0265	0.0375
	NDCG@3	0.0254	0.0337	0.0316	0.0312	0.0380	0.0392	0.0509
	NDCG@5	0.0299	0.0350	0.0335	0.0364	0.0356	0.0373	0.0514
	NDCG@10	0.0315	0.0376	0.0379	0.0373	0.0396	0.0382	0.0518
Epinions	Precision@3	0.0046	0.0138	0.0096	0.0100	0.0105	0.0156	0.0208
	Precision@5	0.0042	0.0104	0.0089	0.0090	0.0098	0.0123	0.0173
	Precision@10	0.0035	0.0080	0.0066	0.0071	0.0086	0.0102	0.0140
	NDCG@3	0.0099	0.0175	0.0136	0.0176	0.0160	0.0183	0.0226
	NDCG@5	0.0128	0.0177	0.0152	0.0196	0.0183	0.0182	0.0217
	NDCG@10	0.0169	0.0202	0.0198	0.0202	0.0200	0.0217	0.0234

Table 2: Performance comparison of different recommender systems




Figure 3: Effect of λ on Ciao dataset.

the performance with varied values of λ using Precision@3 as the measurement. The performance first increases as the value of λ gets larger and then starts to decrease once λ goes beyond 100. The performance weakly depends on the parameter controlling the bidirectional influence, which suggests that transferring user’s information from the social domain to the item domain already significantly boosts the performance. However, the user-item interactions and user-user connections are often very sparse, so the bidirectional mapping (Cycle Reconstruction) is proposed to help alleviate this data sparsity problem. Although the performance weakly depends on the bidirectional influence, we still observe that we can learn better user’s representation in both domains.

4 Related Work

As suggested by the social theories [Marsden and Friedkin, 1993], people’s behaviours tend to be influenced by their social connections and interactions. Many existing social recommendation methods [Fan *et al.*, 2018; Tang *et al.*, 2013a; 2016b; Du *et al.*, 2017; Ma *et al.*, 2008] have shown that incorporating social relations can enhance the performance of the recommendations. In addition, deep neural networks have been adopted to enhance social recommender systems. DLMF [Deng *et al.*, 2017] utilizes deep auto-encoder to initialize vectors for matrix factorization. DeepSoR [Fan *et al.*, 2018] utilizes deep neural networks to capture non-linear user representations in social relations and integrate them into probabilistic matrix factorization for prediction. GraphRec [Fan *et al.*, 2019] proposes a graph neural net-

works framework for social recommendation, which aggregates both user-item interactions information and social interaction information when performing prediction.

Some recent works have investigated adversarial learning for recommendation. IRGAN [Wang *et al.*, 2017] proposes to unify the discriminative model and generative model with adversarial learning strategy for item recommendation. NMRN-GAN [Wang *et al.*, 2018b] introduces the adversarial learning with negative sampling for streaming recommendation. Despite the compelling success achieved by many works, little attention has been paid to social recommendation with adversarial learning. Therefore, we propose a deep adversarial social recommender system to fill this gap.

5 Conclusion and Future Work

In this paper, we present a Deep Adversarial Social recommendation model (**DASO**), which learns separated user representations in item domain and social domain. Particularly, we propose to transfer users’ information from social domain to item domain by using a bidirectional mapping method. In addition, we also introduce the adversarial learning to optimize our entire framework by generating informative negative samples. Comprehensive experiments on two real-world datasets show the effectiveness of our model. The calculation of softmax function in item/social domain generator involves all items/users, which is time-consuming and computationally inefficient. Therefore, hierarchical softmax [Morin and Bengio, 2005; Mikolov *et al.*, 2013; Wang *et al.*, 2018a], which is a replacement for softmax, would be considered to speed up the calculation in both generators in the future direction.

Acknowledgments

The work has been supported, in part, by NSFC-Guangdong Joint Fund under project U1501254, Science Technology and Innovation Committee of Shenzhen Municipality Under project JCYJ20170818095109386, and a start-up fund from the Hong Kong Polytechnic University (project no.1.9B0V). Tyler Derr, Yao Ma and Jiliang Tang are supported by the National Science Foundation (NSF) under grant numbers IIS-1714741, IIS-1715940, IIS-1845081 and CNS-1815636, and a grant from Criteo Faculty Research Award.

References

- [Brock *et al.*, 2019] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019.
- [Cai and Wang, 2018] Liwei Cai and William Yang Wang. Kbgan: Adversarial learning for knowledge graph embeddings. In *NAACL-HLT*, 2018.
- [Deng *et al.*, 2017] Shuiguang Deng, Longtao Huang, Guandong Xu, Xindong Wu, and Zhaohui Wu. On deep learning for trust-aware recommendations in social networks. *TNNLS*, 2017.
- [Derr *et al.*, 2019] Tyler Derr, Hamid Karimi, Xiaorui Liu, Jiejun Xu, and Jiliang Tang. Deep adversarial network alignment. *ArXiv*, abs/1902.10307, 2019.
- [Du *et al.*, 2017] Xixi Du, Huafeng Liu, and Liping Jing. Additive co-clustering with social influence for recommendation. In *RecSys*. ACM, 2017.
- [Fan *et al.*, 2018] Wenqi Fan, Qing Li, and Min Cheng. Deep modeling of social relations for recommendation. In *AAAI*, 2018.
- [Fan *et al.*, 2019] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *WWW*, 2019.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [Jamali and Ester, 2010] Mohsen Jamali and Martin Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *RecSys*. ACM, 2010.
- [Liu *et al.*, 2018] Linqing Liu, Yao Lu, Min Yang, Qiang Qu, Jia Zhu, and Hongyan Li. Generative adversarial network for abstractive text summarization. In *AAAI*, 2018.
- [Ma *et al.*, 2008] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. Sorec: social recommendation using probabilistic matrix factorization. In *CIKM*. ACM, 2008.
- [Mao *et al.*, 2017] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*. IEEE, 2017.
- [Mao *et al.*, 2018] Xudong Mao, Qing Li, Haoran Xie, Raymond Yiu Keung Lau, Zhen Wang, and Stephen Paul Smolley. On the effectiveness of least squares generative adversarial networks. *TPAMI*, 2018.
- [Marsden and Friedkin, 1993] Peter V Marsden and Noah E Friedkin. Network studies of social influence. *Sociological Methods & Research*, 1993.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [Morin and Bengio, 2005] Frederic Morin and Yoshua Bengio. Hierarchical probabilistic neural network language model. In *Aistats*, 2005.
- [Rendle *et al.*, 2009] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*. AUAI Press, 2009.
- [Schulman *et al.*, 2015] John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. In *NIPS*, 2015.
- [Sutton *et al.*, 2000] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, 2000.
- [Tang *et al.*, 2013a] Jiliang Tang, Xia Hu, Huiji Gao, and Huan Liu. Exploiting local and global social context for recommendation. In *IJCAI*, 2013.
- [Tang *et al.*, 2013b] Jiliang Tang, Xia Hu, and Huan Liu. Social recommendation: a review. *Social Network Analysis and Mining*, 2013.
- [Tang *et al.*, 2016a] Jiliang Tang, Charu Aggarwal, and Huan Liu. Recommendations in signed social networks. In *WWW*, 2016.
- [Tang *et al.*, 2016b] Jiliang Tang, Suhang Wang, Xia Hu, Dawei Yin, Yingzhou Bi, Yi Chang, and Huan Liu. Recommendation with social dimensions. In *AAAI*, 2016.
- [Tieleman and Hinton, 2012] T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
- [Wang *et al.*, 2016] Xin Wang, Wei Lu, Martin Ester, Can Wang, and Chun Chen. Social recommendation with strong and weak ties. In *CIKM*. ACM, 2016.
- [Wang *et al.*, 2017] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *SIGIR*. ACM, 2017.
- [Wang *et al.*, 2018a] Hongwei Wang, Jia Wang, Jialin Wang, Miao Zhao, Weinan Zhang, Fuzheng Zhang, Xing Xie, and Minyi Guo. Graphgan: Graph representation learning with generative adversarial nets. In *AAAI*, 2018.
- [Wang *et al.*, 2018b] Qinyong Wang, Hongzhi Yin, Zhiting Hu, Defu Lian, Hao Wang, and Zi Huang. Neural memory streaming recommender networks with adversarial training. In *KDD*. ACM, 2018.
- [Wasserman and Faust, 1994] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*. Cambridge university press, 1994.
- [Zhao *et al.*, 2014] Tong Zhao, Julian McAuley, and Irwin King. Leveraging social connections to improve personalized ranking for collaborative filtering. In *CIKM*. ACM, 2014.