



CDRNP: Cross-Domain Recommendation to Cold-Start Users via Neural Process

Xiaodong Li
Institute of Information Engineering,
Chinese Academy of Sciences.
School of Cyber Security, UCAS[†]
lixiaodong@iie.ac.cn

Jiawei Sheng*
Institute of Information Engineering,
Chinese Academy of Sciences
shengjiawei@iie.ac.cn

Jiangxia Cao
Kuaishou Inc
caojiangxia@kuaishou.com

Wenyuan Zhang
Institute of Information Engineering,
Chinese Academy of Sciences.
School of Cyber Security, UCAS
zhangwenyuan@iie.ac.cn

Quangang Li
Institute of Information Engineering,
Chinese Academy of Sciences
liquangang@iie.ac.cn

Tingwen Liu*
Institute of Information Engineering,
Chinese Academy of Sciences.
School of Cyber Security, UCAS
liutingwen@iie.ac.cn

ABSTRACT

Cross-domain recommendation (CDR) has been proven as a promising way to tackle the user cold-start problem, which aims to make recommendations for users in the target domain by transferring the user preference derived from the source domain. Traditional CDR studies follow the embedding and mapping (EMCDR) paradigm, which transfers user representations from the source to target domain by learning a user-shared mapping function, neglecting the user-specific preference. Recent CDR studies attempt to learn user-specific mapping functions in meta-learning paradigm, which regards each user's CDR as an individual task, but neglects the preference correlations among users, limiting the beneficial information for user representations. Moreover, both of the paradigms neglect the explicit user-item interactions from both domains during the mapping process. To address the above issues, this paper proposes a novel CDR framework with neural process (NP), termed as CDRNP. Particularly, it develops the meta-learning paradigm to leverage user-specific preference, and further introduces a stochastic process by NP to capture the preference correlations among the overlapping and cold-start users, thus generating more powerful mapping functions by mapping the user-specific preference and common preference correlations to a predictive probability distribution. In addition, we also introduce a preference remainder to enhance the common preference from the overlapping users, and finally devises an adaptive conditional decoder with preference modulation to make prediction for cold-start users with items in the target domain. Experimental results demonstrate that CDRNP outperforms previous SOTA methods in three real-world CDR scenarios.

*Corresponding author.

[†]University of Chinese Academy of Sciences.

CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Neural networks.

KEYWORDS

Cross-Domain Recommendation; Cold-Start Recommendation; Meta Learning; Neural Process

ACM Reference Format:

Xiaodong Li, Jiawei Sheng, Jiangxia Cao, Wenyuan Zhang, Quangang Li, Tingwen Liu. 2024. CDRNP: Cross-Domain Recommendation to Cold-Start Users via Neural Process. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3616855.3635794>

1 INTRODUCTION

Recommendation systems (RS) have been widely deployed in real-world applications, such as Amazon (E-commerce) and YouTube (online video). As a promising method, collaborative filtering (CF) has been successfully applied to RS, ranging from early matrix factorization methods [26, 31] to the recent deep neural network methods [9, 35]. However, CF-based methods usually suffer from the long-standing challenge of *user cold-start problem* [39], where the new emerging user has few *user-item interactions* that can hardly learn effective representations to make recommendation. To address this problem, **cross-domain recommendation (CDR)** [15, 18] has been proposed recently, which explores the cold-start users by utilizing the rich user-item interactions observed from an informative *source domain*, so as to improve the cold-start users' recommendation in the *target domain*.

A common idea of CDR methods is to train the model with the *overlapping users* appearing in both the source and target domain, and then make recommendation for the *cold-start users* that only have interactions in the source domain. To achieve the above idea, traditional CDR methods [10, 20, 27, 30, 39] follow the Embedding and Mapping (EMCDR) paradigm (as shown in Fig. 1(a)). This paradigm aims to learn a user-shared mapping function. For implementation, EMCDR first generates user/item representations for each domain separately, and then trains a mapping function to align the overlapping users representations between the source and



This work is licensed under a Creative Commons Attribution International 4.0 License.

WSDM '24, March 4–8, 2024, Merida, Mexico

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0371-3/24/03.

<https://doi.org/10.1145/3616855.3635794>

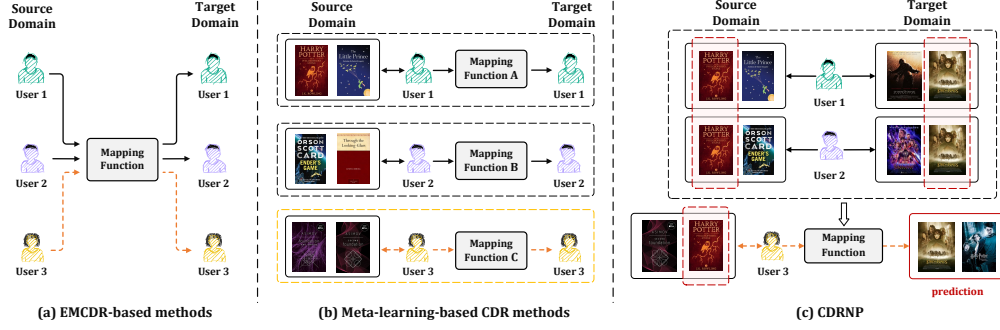


Figure 1: Users 1 and 2 are overlapping users, while User 3 is a cold-start user. (a) The framework of EMCDR-based methods by learning a user-shared mapping function. (b) The framework of meta-learning-based CDR methods by learning user-specific mapping functions. (c) The framework of our CDRNP by capturing preference correlations among users.

target domain. Afterward, given a cold-start user representation from the source domain, EMCDR uses the mapping function to predict the user representations in the target domain. However, such a paradigm learns a unified mapping function shared by all users, neglecting the different user-specific preferences, thus leading to inflexible representation transferring for different users.

Recent studies attempt meta-learning [6, 22] to further capture user-specific preference in CDR (as shown in Fig. 1(b)). They mostly regard each user’s CDR as an individual task, and train the model to achieve these user-specific tasks with the users in the support set, thus can make recommendation for cold-start users in the query set. Among them, PTUPCDR [40] proposes a meta network with a task-oriented optimization procedure, and generates cold-start user representation by personalized mapping functions in prediction. However, this kind of methods usually regard each user’s CDR as an independent task, neglecting the preference correlations among users, where the similar users may have similar preferences which can provide beneficial information for training the mapping functions. Moreover, both of the EMCDR-based and the meta-learning-based methods mainly concentrate on the transferring of the learned user representations, neglecting the explicit user-item interactions from both domains during the mapping process, thus leading to sub-optimal user preference modeling.

To address the above issues, we propose a novel meta-learning framework for CDR based on Neural Process [7, 8], termed as CDRNP (as shown in Fig. 1(c)). Specifically, we build our model with Neural Process (NP) principle, which formulates the CDR task of all users within the same stochastic process, naturally capturing the preference correlations among the overlapping and cold-start users. Besides, the model also follows the meta-learning paradigm, which not only preserves user-specific information for each user, but also utilizes explicit user-item interactions from both domains beyond the sole user representation transferring. Such a design enables CDRNP to learn more representative user-specific mapping functions and achieve bi-directional recommendation with same observed data. For implementation of the CDRNP framework, we first propose a *characteristic embedding layer* which leverages each user’s interactions observed in the source domain to generate its user-specific preference representation, and then combine it with the candidate item in the target domain, constituting the complete

rating-observed information as input for each candidate item rating. Then, we propose a *preference aggregator* to approximate CDR for each user as a particular instance of a stochastic process, which captures common preference correlations among the cold-start users and overlapping users. Additionally, we devise a *preference remainder* to enhance the common preference representations from the overlapping users. Finally, we devise an *adaptive conditional decoder* to modulate both preference representations to make better rating prediction for the cold-start users in the target domain.

In summary, the contributions of this paper are as follows:

- To our best knowledge, we introduce a fresh perspective to solve the cold-start problem in CDR with NP principle.
- We propose a novel meta-learning CDR method with NP, termed as CDRNP, which preserves the capability of user-specific mapping function by meta-learning paradigm, while further captures preference correlations among users with NP.
- Extensive experiments on three CDR scenarios demonstrate that CDRNP outperforms the state-of-the-art methods.

2 PRELIMINARY

This section introduces the CDR problem settings and the background knowledge of the neural process principle.

2.1 Problem Setting

This paper considers a general CDR scenario: given the observed user-item interactions of both domains, we aim to make recommendations in the target domain for cold-start users, with the help of user-item interactions in the source domain. Formally, the data can be formulated as $\mathcal{D}^{s/t} = (\mathcal{U}^{s/t}, \mathcal{V}^{s/t}, \mathcal{Y}^{s/t})$, where the superscript s/t denotes the source/target domain, and \mathcal{U} , \mathcal{V} and $\mathcal{Y} \in \{0, 1, 2, 3, 4, 5\}^{|\mathcal{U}| \times |\mathcal{V}|}$ denote the user set, item set and user-item rating matrix, respectively. Here the rating score $y(u, v) \in \mathcal{Y}$ ranges from 0 to 5, reflecting the preference tendency of user $u \in \mathcal{U}$ to choose item $v \in \mathcal{V}$. Further, we define the users overlapped in both domains as the overlapping users $\mathcal{U}^o = \mathcal{U}^s \cap \mathcal{U}^t$, then the cold-start users from the source domain can be denoted as $\mathcal{U}^{s \setminus o} = \mathcal{U}^s \setminus \mathcal{U}^o$. Therefore, for each cold-start user $u^{s \setminus o} \in \mathcal{U}^{s \setminus o}$, given its interactions from the source domain $\mathcal{V}^s(u^{s \setminus o})$, the goal is to predict the rating $y(u^{s \setminus o}, v^t)$ for each candidate $v^t \in \mathcal{V}^t$.

To achieve the above goal, we develop our model with meta-learning settings [1, 3, 13, 19, 33] for users: For each user $u_j \in \mathcal{U}$, to predict rating $y(u_j, v_k^t)$ in the target domain, we briefly summarize the rating-observed information as $x_{jk} = (u_j, \mathcal{V}^s(u_j), v_k^t)$. Here we use $x_m \triangleq x_{jk}$ and $y_m^t \triangleq y(u_j, v_k^t)$ for brevity. Based upon, we define the concept of task ω_i in meta-learning, where each $\omega_i = \{(x_m, y_m^t) | (x_m, y_m^t) \in \{C_i \cup Q_i\}\}$ batched with multiple users. Here, we divide ω_i into the support set C_i and the query set Q_i , such that $\omega_i = C_i \cup Q_i$.

Note that we train our model on the tasks $\omega_i \in \Omega_{\text{meta-training}}$ consisting of the overlapping users, i.e., we use (x_m^o, y_m^t) pairs with $u_j^o \in \mathcal{U}^o$ in both C_i and Q_i , so as to learn the common knowledge of achieving CDR with these tasks. In the testing phase, our model predicts the interactions of the cold-start users in target domain via new tasks $\omega_i \in \Omega_{\text{meta-testing}}$, where we use (x_m^o, y_m^t) pairs with $u_j^o \in \mathcal{U}^o$ in C_i and we predict $(x_m^{s\setminus o}, y_m^t)$ pairs with $u_j^{s\setminus o} \in \mathcal{U}^{s\setminus o}$ in Q_i . In such a way, the model learns common knowledge of achieving CDR in the training phase, thus can make recommendation for the cold-start users to achieve CDR in the testing phase.

2.2 Neural Process Principle

To achieve CDR, we develop our model based on Neural Process (NP) [8]. Generally, NP is a kind of neural latent variable models, which approximates stochastic processes parameterized by powerful neural networks. In this paper, we assume each task $\omega_i = \{(x_m, y_m^t) | (x_m, y_m^t) \in \{C_i \cup Q_i\}\}$ associated with an instantiation of the stochastic process f_i , and each pair (x_m, y_m^t) can be seen as a sample point drawn from f_i . Therefore, the corresponding probability distribution ρ of f_i can be defined as follows:

$$\rho_{x_{1:M_i}}(y_{1:M_i}^t) = \int p(f_i) p(y_{1:M_i}^t | f_i, x_{1:M_i}) df_i, \quad (1)$$

where $x_{1:M_i}$ and $y_{1:M_i}^t$ denote the all x_m and y_m^t decoupled from ω_i , respectively, and we use $M_i = |\omega_i|$ for simplicity. Following NP, we approximate the stochastic process f_i with a low-dimensional random vector z_i generated by non-linear neural networks. Then, with *i.i.d.* condition on tasks, we rewrite the generation process as:

$$p(y_{1:M_i}^t | x_{1:M_i}) = \int p(z_i) \prod_{m=1}^{M_i} p(y_m^t | x_m, z_i) dz_i, \quad (2)$$

Here we assume $p(z_i) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ for simplicity in paradigmatic variational frameworks [12], and then we can easily sample z_i from $p(z_i)$ to achieve concrete function implementation.

Since the true posterior distribution is intractable, we adopt amortised variational inference [21] to learn it. Let $q(z_i | \omega_i)$ be the variational posterior of the latent variable z_i , the corresponding evidence lower-bound (ELBO) derivation is as follows:

$$\begin{aligned} \log p(y_{1:M_i}^t | x_{1:M_i}) &= \log \int p(z_i, y_{1:M_i}^t | x_{1:M_i}) dz_i, \\ &\geq E_{q(z_i | \omega_i)} \left[\log \frac{p(z_i, y_{1:M_i}^t | x_{1:M_i})}{q(z_i | \omega_i)} \right], \\ &= E_{q(z_i | \omega_i)} \left[\sum_{m=1}^{M_i} \log p(y_m^t | x_m, z_i) + \log \frac{p(z_i)}{q(z_i | \omega_i)} \right], \end{aligned} \quad (3)$$

Considering that each task ω_i contains a support set C_i and a query set Q_i , we would like to make predictions for Q_i given C_i . Therefore,

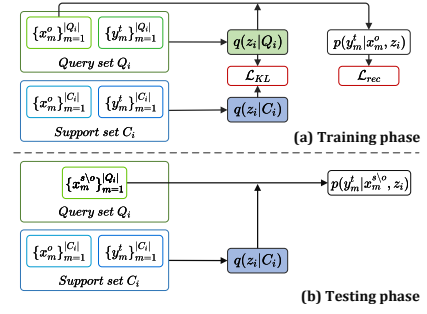


Figure 2: An illustration of our model CDRNP with NP in the training and testing phase.

we rewrite Eq. (3) as follows:

$$\log p(y_{1:|Q_i|}^t | x_{1:|Q_i|}, C_i) \quad (4)$$

$$\geq E_{q(z_i | Q_i)} \left[\sum_{m=1}^{|Q_i|} \log p(y_m^t | x_m, z_i) + \log \frac{p(z_i | C_i)}{q(z_i | Q_i)} \right], \quad (5)$$

$$= \underbrace{-E_{q(z_i | Q_i)} \log p(y_{1:|Q_i|}^t | x_{1:|Q_i|}, z_i)}_{\text{desired log-likelihood}} + \underbrace{\text{KL}(q(z_i | Q_i) || q(z_i | C_i))}_{\text{KL divergence}} \quad (6)$$

where Eq. (5) is the lower bound of Eq. (4) with derivation like Eq. (3). In practice, $p(z_i | C_i)$ in Eq. (5) is intractable, thus we use variational prior $q(z_i | C_i)$ to approximate $p(z_i | C_i)$. Such an approximation can be regarded as a regularization term of KL divergence to ensure the *consistency* condition of stochastic process [24]. Thereby, we summarize Eq. (5) into Eq. (6), which can be intuitively explained with two terms: the first term aims to achieve the desired CDR task, and the second term aims to ensure all tasks in the query and support set following the same stochastic process.

For implementation, we illustrate the above principle from a practical perspective. Fig. 2 shows our model with NP in the training and testing phase. Note that in the training phase, we use \mathcal{L}_{KL} to assimilate $q(z_i | Q_i)$ and $q(z_i | C_i)$, which learns z_i general to both the support and query set in each task. In the testing phase, we use z_i generated from the $q(z_i | C_i)$ to substitute for $q(z_i | Q_i)$, and thus together with $\{x_m^{s\setminus o}\}_{m=1}^{|Q_i|}$ in the given task, we make predictions for the users by $p(y | x, z)$. In this way¹, we encourage the tasks of users following the same stochastic process, thus naturally captures the preference correlations among overlapping and cold-start users.

3 MODEL

This section introduces our model CDRNP implementing the NP principle. Fig. 3 sketches the overall framework of CDRNP. Our model consists of three components: 1) A characteristic embedding layer, which initializes user/item embeddings and generates the rating-observed information representations for users. 2) A preference encoder, which contains a preference aggregator to achieve NP of the support and query set, and a preference remainder to enhance common preference representations from the support set. 3)

¹Note that in the training phase, we leverage overlapping users u_j^o in both Q_i and C_i , so we can make prediction on Q_i by imitating C_i . In the testing phase, we make predictions for cold-start users $u_j^{s\setminus o}$ in Q_i , by imitating overlapping users u_j^o in C_i .




Figure 3: The framework of CDRNP in the training phase. The characteristic embedding layer is used to generate x_m^o . Both of the C_i and Q_i are encoded to generate the variational prior and posterior. h_i learned from the preference remainder is used to modulate the decoder parameters. z_i sampled from $q(z_i|Q_i)$ is concatenated with x_m^o in Q_i to predict y_m^t for Q_i via decoder.

An adaptive conditional decoder to incorporate features from the preference aggregator and preference remainder, and finally predict ratings of items in the target domain for the cold-start users.

3.1 Characteristic Embedding Layer

In CDR, our goal is to predict y_m^t in the (x_m^o, y_m^t) pair for the cold-start users, given the (x_m^o, y_m^t) pairs of overlapping users. Here x_m (with superscript $s \setminus o$ or o) summarizes the rating-observed information $x_m \triangleq x_{jk} = (u_j, \mathcal{V}^s(u_j), v_k^t)$ for user u_j to choose v_k^t with the user's interactions $\mathcal{V}^s(u_j)$ in the source domain. To achieve the above goal, we initialize the user/item representations in both domains, and then propose an attentive characteristic preference module to adaptively capture user's interactions $\mathcal{V}^s(u_j)$ in the source domain. Finally, we fuse the user's information from the source domain with the candidate item v_k^t in the target domain, so as to complete the rating-observed information x_m .

Specifically, we first initialize user embeddings as $u_j \in \mathbb{R}^d$, and items embeddings as $v_k^s \in \mathbb{R}^{d^2}$ in the source domain and $v_k^t \in \mathbb{R}^{d^2}$ in the target domain². Intuitively, various items that users interact in the source domain contribute differently to their preferences. Therefore, we propose an attentive characteristic preference module over these item embeddings with attention mechanism [36]:

$$\begin{aligned} a_k &= q^T \text{ReLU}(W_a v_k^s), \\ a'_k &= \frac{\exp(a_k)}{\sum_{v_n^s \in \mathcal{V}^s(u_j)} \exp(a_n)}, \\ c_{u_j} &= \sum_{v_k^s \in \mathcal{V}^s(u_j)} a'_k v_k^s, \end{aligned} \quad (7)$$

where $q \in \mathbb{R}^d$ and $W_a \in \mathbb{R}^{d \times d}$ are learnable parameters. a'_k is the attention score of item v_k^s , differentiating the contribution of item v_k^s to the user preference. Thereby, $c_{u_j} \in \mathbb{R}^{d^2}$ summarizes the user-specific preference derived from the source domain.

²Here d is the dimension of user embedding. As suggested in previous work [40], we let the dimension of item embedding be d^2 for the convenience of Eq. (8).

To predict rating y_m^t of candidate item v_k^t for each user, we enhance the user feature (including u_j and c_{u_j}) with the item embedding v_k^t in the target domain. Specifically, we first fuse u_j and c_{u_j} to generate the user feature derived from the source domain, and then associate it with the candidate item:

$$\hat{u}_j = \text{Resize}(c_{u_j})u_j, \quad x_m = [\hat{u}_j || v_k^t], \quad (8)$$

where we resize $c_{u_j} \in \mathbb{R}^{d^2}$ into $\mathbb{R}^{d \times d}$, which serves as a user-specific projection W_{u_j} in linear transformation, allowing better user feature fusion [40]. Besides, $[\cdot || \cdot]$ denotes concatenation operation. Finally, we output $x_m \in \mathbb{R}^d$ as the rating-observed information representation for the user u_j with candidate item v_k^t .

Note that in the training phase, we leverage overlapping users in both Q_i and C_i , so we can learn how to make prediction on Q_i by imitating C_i . In the testing phase, we make predictions for cold-start users in Q_i , by imitating the overlapping users in C_i .

3.2 Preference Encoder

This section introduces two modules: 1) a preference aggregator to capture preference correlations among users via NP; 2) a preference remainder to enhance original user preference features.

3.2.1 Preference Aggregator. Recall that, we regard each task $\omega_i = \{(x_m, y_m^t) | (x_m, y_m^t) \in \{C_i \cup Q_i\}\}$ decoupled from the same stochastic process f_i , and we approximate f_i by representation z_i generated from neural networks. Therefore, given Q_i and C_i , we aim to generate the variational approximations $q(z_i|Q_i)$ and $q(z_i|C_i)$.

As suggested in previous NP studies [8], we incorporate the (x_m, y_m^t) pair to approximate the stochastic process with neural networks, which can be formulated as:

$$r_{i,m} = \text{MLP}([x_m || y_m^t]; \theta), \quad (x_m, y_m^t) \in \omega_i \quad (9)$$

where $x_m \in \mathbb{R}^d$, and $y_m^t \in \mathbb{R}^1$ denotes a vector with one dimension of the scalar y_m^t . Here we use $\text{MLP}(\cdot; \theta)$ with three-layer neural networks parameterized by θ to fuse the features of x_m and y_m^t .

Afterward, to obtain task representation \bar{r}_i , we aggregate $r_{i,m}$ with a commutative operation as in NP studies [8]. Taking $q(z_i|C_i)$

as an example, the aggregation process is as follows:

$$\begin{aligned}\bar{\mathbf{r}}_i &= \mathbf{r}_{i,1} \oplus \mathbf{r}_{i,2} \oplus \mathbf{r}_{i,3} \oplus \cdots \oplus \mathbf{r}_{i,|C_i|-1} \oplus \mathbf{r}_{i,|C_i|}, \\ &= \frac{1}{|C_i|} \sum_{m=1}^{|C_i|} \mathbf{r}_{i,m},\end{aligned}\quad (10)$$

where \oplus denotes a commutative operation, achieved by vector addition operation in our model for computational efficiency. Note that $\bar{\mathbf{r}}_i$ is a permutation-invariance representation, which is a crucial property to ensure the *exchangeability* condition of stochastic process [24]. Such a design reduces the time complexity to $O(n+m)$, where n and m are the number of users in C_i and Q_i , respectively.

To derive the random variable \mathbf{z}_i of the stochastic process, as widely studied in variational frameworks [12], we employ Gaussian distribution to represent \mathbf{z}_i , i.e., imposing $\mathbf{z}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\sigma}_i^2)$. Therefore, we generate \mathbf{z}_i from the task representation $\bar{\mathbf{r}}_i$ as:

$$\begin{aligned}\hat{\mathbf{r}}_i &= \text{ReLU}(\mathbf{W}_r \bar{\mathbf{r}}_i), \\ \boldsymbol{\mu}_i &= \mathbf{W}_\mu \hat{\mathbf{r}}_i, \quad \log \boldsymbol{\sigma}_i = \mathbf{W}_\sigma \hat{\mathbf{r}}_i,\end{aligned}\quad (11)$$

where $\mathbf{W}_r \in \mathbb{R}^{d \times d}$, $\mathbf{W}_\mu \in \mathbb{R}^{d \times d}$, $\mathbf{W}_\sigma \in \mathbb{R}^{d \times d}$ are learnable parameters. However, the random variable is intractable in training neural networks, we employ the well-known reparameterization trick [12] to sample data from probability distribution, thus the gradients can be smoothly backpropagated:

$$\mathbf{z}_i = \boldsymbol{\mu}_i + \boldsymbol{\epsilon} \odot \boldsymbol{\sigma}_i, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}), \quad (12)$$

where \odot denotes the element-wise product operation, $\boldsymbol{\epsilon}$ is the Gaussian noise. In this way, we can sample \mathbf{z}_i from both $q(\mathbf{z}_i|Q_i)$ and $q(\mathbf{z}_i|C_i)$ in Q_i and C_i , respectively. Besides, by formulating all tasks involving users' interactions into the same stochastic process, we capture preference correlations among users for the CDR problem.

3.2.2 Preference Remainer. To further enhance the preference correlations among users in different layer-depths, we devise a preference remainder to retain the original overlapping users preference features from the support set C_i . Specifically, we first generate user preference representation $\mathbf{h}_{i,m}$ by fusing rating-observed information \mathbf{x}_m^o and the item rating y_m^t in the target domain, observed in the given support set C_i as:

$$\mathbf{h}_{i,m} = \text{MLP}([\mathbf{x}_m^o || \mathbf{y}_m^t]; \phi), \quad (\mathbf{x}_m^o, y_m^t) \in C_i \quad (13)$$

where $\text{MLP}(\cdot; \phi)$ denotes three-layer neural networks parameterized by ϕ and $[\cdot || \cdot]$ denotes the concatenation operation. Then, we derive the common preference representation \mathbf{h}_i by aggregating each user preference representation $\mathbf{h}_{i,m}$ as follows:

$$\begin{aligned}\mathbf{h}_i &= \mathbf{h}_{i,1} \oplus \mathbf{h}_{i,2} \oplus \mathbf{h}_{i,3} \oplus \cdots \oplus \mathbf{h}_{i,|C_i|-1} \oplus \mathbf{h}_{i,|C_i|}, \\ &= \frac{1}{|C_i|} \sum_{m=1}^{|C_i|} \mathbf{h}_{i,m},\end{aligned}\quad (14)$$

Here we use the same operation in Eq. (10) for convenience. In this way, we derive the common preference representations from the support set C_i , enriching the preference correlations among users for making prediction in the query set Q_i .

3.3 Adaptive Conditional Decoder

To make prediction for rating y_m^t given \mathbf{x}_m in Q_i , we have obtained the preference information from C_i , including \mathbf{z}_i from the preference aggregator and \mathbf{h}_i from the preference remainder. To

better incorporate \mathbf{z}_i and \mathbf{h}_i for predicting item ratings, we propose an adaptive conditional decoder based on the well-designed FiLM [5, 25], which can fully modulate the two features in different layer-depths for decoding.

Specifically, the adaptive conditional decoder learns the conditional likelihood $p(y_m^t | \mathbf{x}_m, \mathbf{z}_i)$ (in Eq.(5)) with the \mathbf{x}_m and \mathbf{z}_i (in Eq. (12)) as inputs. We first generate two modulation coefficients $\boldsymbol{\gamma}_i$ and $\boldsymbol{\beta}_i$ over the representation \mathbf{h}_i , and then the adaptive conditional decoder can be formulated as:

$$\begin{aligned}\boldsymbol{\gamma}_i^l &= \tanh(\mathbf{h}_i; \mathbf{W}_\gamma^l), \quad \boldsymbol{\beta}_i^l = \tanh(\mathbf{h}_i; \mathbf{W}_\beta^l), \\ \mathbf{g}^0 &= [\mathbf{x}_m || \mathbf{z}_i], \quad \mathbf{g}^{l+1} = \text{ReLU}(\boldsymbol{\gamma}_i^l \odot (\mathbf{W}_g^l \mathbf{g}^l + \mathbf{b}_g^l) + \boldsymbol{\beta}_i^l),\end{aligned}\quad (15)$$

where \mathbf{g}^l is the input of l -th layer and $[\cdot || \cdot]$ denotes the concatenation operation. We use \tanh to enforce the modulation coefficients in range $[-1, 1]$. In this way, $\boldsymbol{\gamma}_i^l$ controls the gain from \mathbf{g}^l , while $\boldsymbol{\beta}_i^l$ remains the features of \mathbf{h}_i . By stacking multiple layers, the adaptive conditional decoder fully captures preference correlations among overlapping and cold-start users in different layer-depths, and finally output the rating \hat{y}_m^t of user u_j on item v_k^t at the last layer.

3.4 Loss Function

To train our model, the loss function consists of two terms: the desired log-likelihood to achieve CDR, and the KL divergence to impose regularization. Particularly, we rephrase the likelihood \mathcal{L}_{rec} in Eq. (6) as a regression-based loss function by MSE loss:

$$\begin{aligned}\mathcal{L}_{rec} &= -E_{q(\mathbf{z}_i|Q_i)} \log p(y_{1:|Q_i|}^t | \mathbf{x}_{1:|Q_i|}^o, \mathbf{z}_i), \\ &\propto \frac{1}{|Q_i|} \sum_{m=1}^{|Q_i|} (y_m^t - \hat{y}_m^t)^2,\end{aligned}\quad (16)$$

where we use overlapping users $\mathbf{x}_{1:|Q_i|}^o$ to replace $\mathbf{x}_{1:|Q_i|}$. Besides, we impose \mathcal{L}_{KL} for the *consistency* condition in the approximated stochastic process, which practically minimizes the KL divergence in Eq. (6) between the approximate posterior of the query set Q_i and conditional prior of the support set C_i :

$$\mathcal{L}_{KL} = \text{KL}(q(\mathbf{z}_i|Q_i) || q(\mathbf{z}_i|C_i)), \quad (17)$$

In summary, the overall loss function is defined as follows:

$$\mathcal{L} = \mathcal{L}_{rec} + \lambda \mathcal{L}_{KL}, \quad (18)$$

where $\lambda \in [0, 1]$ is the hyper-parameter to balance losses.

Note that we use overlapping users in the training phase, CDRNP first learns on support set C_i and then updates according to the prediction loss over query set Q_i . In the testing phase, we randomly removed all the ground-truth rating y_m^t in Q_i and regard them as the test users (cold-start users). CDRNP can predict the preference of the cold-start users in Q_i via a new task ω_i with only overlapping users in C_i available and the global information learned from the training phase. Thus given Q_i and C_i , the decoder takes \mathbf{z}_i sampled from $q(\mathbf{z}_i|C_i)$ and $\mathbf{x}_m^{s|o}$ in Q_i as inputs to predict \hat{y}_m^t for Q_i .

4 EXPERIMENTS

In this section, we compare our proposed model with previous models, and conduct experiments to evaluate our model.

Table 1: Statistics of Three CDR scenarios (*Overlap* denotes the number of overlapping users).

Scenarios	#Items		#Users			#Ratings	
	Source	Target	Overlap	Source	Target	Source	Target
Movie → Music	50,052	64,443	18,031	123,960	75,258	1,697,533	1,097,592
Book → Movie	367,982	50,052	37,388	603,668	123,960	8,898,041	1,697,533
Book → Music	367,982	64,443	16,738	603,668	75,258	8,898,041	1,097,592

4.1 Datasets

Following existing methods [10, 37, 39] on CDR, we evaluate CDRNP on a real-world recommendation dataset, namely the Amazon review dataset. Specifically, we use the Amazon-5scores dataset where each user or item has at least five ratings.

Following [10, 37], we select three popular categories out of 24 categories for CDRNP: movies_and_tv (Movie), cds_and_vinyl (Music), and books (Book). We define three CDR scenarios: Movie-domain→Music-domain, Book-domain→Movie-domain, and Book-domain→Music-domain. While many existing works only evaluate a subset of data, we conduct experiments using all available data, which more realistically simulates the real-world. Table 1 summarizes the detailed information of three CDR scenarios.

4.2 Experimental Setting

4.2.1 Evaluation Metrics. Following previous work [20, 37], we use mean absolute error (MAE) and root mean square error (RMSE) as evaluation metrics to assess the recommendation performance. For each metric, we report the average results over five random runs for all users. Following previous meta-learning recommenders [1, 3, 13], we only predict the score of each user on candidate items in the query set \mathcal{Q}_i .

4.2.2 Compared Baselines. For a fair comparison, we choose the following methods as baselines: (1) **TGT** represents the target MF model, which only uses data from the target domain for training and ignores user interactions in the source domain. (2) **CMF** [31] leverages the source domain interactions as additional features for overlapping users to achieve performance improvement in the target domain. (3) **EMCDR** [20] follows the embedding and mapping idea for CDR, learning specific user representations in each domain through a mapping function, which is then utilized to transfer user representations from the source domain to the target domain. (4) **DCDCSR** [38] is a CDR model based on matrix factorization and DNNs. It generates latent vectors for users and items using MF model, then uses DNNs to map latent vectors across domains. (5) **SSCDR** [10] proposes a CDR framework based on semi-supervised mapping. It trains the mapping function by exploiting overlapping users and all items. Moreover, it predicts the latent vectors of cold-start users by aggregating neighborhood information. (6) **PTUPCDR** [40] is a meta-learning-based CDR method that generates personalized mapping functions for each user by learning a meta-network, transferring the user’s preference from the source domain to the target domain.

4.2.3 Implementation and Hyperparameter Setting. The source codes for all baselines have been released, and we modify the data input and evaluation sections to fit our experimental settings. To ensure

a fair comparison, we set the same hyper-parameter settings for all methods. For each method, we tune the dimension of embedding in {32, 64, 128, 256}. All models are trained for 10 epochs to achieve convergence. We use the same fully connected layer to facilitate comparison for CDRNP and other baselines. The stacking layers for preference aggregator, preference remainder and adaptive conditional decoder in our model are three, for attentive characteristic preference module is two. The learning rate is chosen from {0.001, 0.005, 0.01, 0.02, 0.1}. The hidden layer size for each layer is chosen from {8, 16, 32, 64, 128}, and the hyper-parameter λ in Eq. (18) is chosen from {0.1, 0.3, 0.5, 0.7, 0.9}. In our experiments, we set the proportion of the testing dataset users α to be 20%, 50%, and 80% of the total overlapping users, respectively. We consider the sequential timestamps to avoid information leakage.

4.3 Performance Comparisons on CDR

This section presents the experimental results and detailed analyses of CDRNP in three CDR scenarios and in-depth discussion of the CDRNP bi-directional recommendation performance for cold-start users in both the source and target domains.

4.3.1 Performance Comparisons. Table 2 shows the experimental performance of all methods in three CDR scenarios. Note that the size of \mathcal{C}_i is set to 40, and the length of interaction items history $\mathcal{V}^s(u_j)$ is set to 20. Based on the experimental results, we can draw the following conclusions: (1) *CDRNP consistently exhibits the best performance in all CDR scenarios.* Specifically, under the experimental setting of α value at 20%, CDRNP outperforms the existing SOTA (i.e., PTUPCDR) in MAE by 30.68%, 11.27% and 39.34% in three CDR scenarios. (2) *CDRNP demonstrates outstanding performance and robustness even when the training dataset is limited.* For example, for Movie domain→Music domain, with α value set to 20%, 50%, and 80%, CDRNP achieves improvements of 30.68%, 37.76%, and 41.06%, respectively, compared to the state-of-the-art meta-learning method. This indicates that our NP framework can effectively adapt to CDR even with limited trainable overlapping users. (3) *CDRNP achieves competitive performance compared with all EMCDR-based and meta-learning-based baselines.* The results highlight the necessity of capturing the preference correlations among overlapping and cold-start users.

4.3.2 Bi-directional Recommendation Performance. We evaluate the bi-directional recommendation performance of CDRNP to cold-start users in both the source and target domain with EMCDR and PTUPCDR. Specifically, we first split 50% of the overlapping users for training EMCDR and PTUPCDR in the Movie-domain→Music-domain. To ensure fairness, we split 25% of overlapping users for testing EMCDR and PTUPCDR in Movie-domain→Music-domain and the remained 25% of overlapping users for testing in Music-domain→Movie-domain, respectively. We take the same operation to train CDRNP, EMCDR and PTUPCDR in the Music-domain→Movie-domain, where testing them in the Music-domain→Movie-domain and Movie-domain→Music-domain. The empirical results are presented in Figure 3, from which we can draw the following conclusions: (1) *CDRNP achieved the best bi-directional recommendation performance to cold-start users in both the source and target domains.*

Table 2: Performance comparison of CDR to cold-start users in three CDR scenarios.

Methods	Movie-domain→Music-domain						Book-domain→Movie-domain						Book-domain→Music-domain					
	20%		50%		80%		20%		50%		80%		20%		50%		80%	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
TGT	4.4803	5.1580	4.4989	5.1736	4.5020	5.1891	4.1831	4.7536	4.2288	4.7920	4.2123	4.8149	4.4873	5.1672	4.5073	5.1727	4.5204	5.2308
CMF	1.5209	2.0158	1.6893	2.2271	2.4186	3.0936	1.3632	1.7918	1.5813	2.0886	2.1577	2.6777	1.8284	2.3829	2.1282	2.7275	3.0130	3.6948
DCDCSR	1.4918	1.9210	1.8144	2.3439	2.7194	3.3065	1.3971	1.7346	1.6731	2.0551	2.3618	2.7702	1.8411	2.2955	2.1736	2.6771	3.1405	3.5842
SSCDR	1.3017	1.6579	1.3762	1.7477	1.5046	1.9229	1.2390	1.6526	1.2137	1.5602	1.3172	1.7024	1.5414	1.9283	1.4739	1.8441	1.6414	2.1403
EMCDR	1.2350	1.5515	1.3277	1.6644	1.5008	1.8771	1.1162	1.4120	1.1832	1.4981	1.3156	1.6433	1.3524	1.6737	1.4723	1.8000	1.7191	2.1119
PTUPCDR	1.1504	1.5195	1.2804	1.6380	1.4049	1.8234	0.9970	1.3317	1.0894	1.4395	1.1999	1.5916	1.2286	1.6085	1.3764	1.7447	1.5784	2.0510
CDRNP	0.7974	1.0638	0.7969	1.0589	0.8280	1.0758	0.8846	1.1327	0.8946	1.1450	0.8970	1.1576	0.7453	0.9914	0.7629	1.0111	0.7787	1.0373
Improve	30.68%	29.99%	37.76%	35.35%	41.06%	41.00%	11.27%	14.94%	17.88%	20.46%	25.24%	27.27%	39.34%	38.36%	44.57%	42.05%	50.67%	49.42%

Table 3: Performance of bi-directional recommendation.

Model	Movie → Music		Music → Movie	
	MAE	RMSE	MAE	RMSE
EMCDR _{Movie→Music}	1.6573	2.0170	4.1198	4.5604
EMCDR _{Music→Movie}	4.3590	4.9620	1.1773	1.5165
PTUPCDR _{Movie→Music}	1.2859	1.7384	4.4992	5.3533
PTUPCDR _{Music→Movie}	4.3781	4.9627	1.0954	1.4161
CDRNP	0.9140	1.1257	0.8817	1.1383



Table 4: Ablation Study on three CDR scenarios.

Model	Metric	Scenario 1	Scenario 2	Scenario 3
CDRNP(w/o PRM+ACP)	MAE	0.8597	0.9446	0.7940
	RMSE	1.0908	1.1645	1.0214
CDRNP(w/o PRM)	MAE	0.8259	0.9002	0.7628
	RMSE	1.0735	1.1409	1.0019
CDRNP	MAE	0.7974	0.8846	0.7453
	RMSE	1.0638	1.1327	0.9914

CDRNP only needs to be trained in the Music-domain→Movie-domain to achieve outstanding bi-directional recommendation performance for cold-start users in both the Music-domain and Movie-domain. (2) *In comparison with EMCDR and PTUPCDR, CDRNP can achieve outstanding CDR performance to cold-start users in the target domain.* This highlights the necessity of utilizing explicit user-item interactions from both domains.

4.4 Ablation Study

In this section, we conduct an ablation study to test the effectiveness of each module in CDRNP. We first remove the Preference Remainer (PRM) module (FiLM layer number 0) from CDRNP, and then remove the Attentive Characteristic Preference (ACP) module based on that. We evaluate the intact CDRNP (FiLM layer number 3). We utilize the α value of 20% for all CDR scenarios. From the results shown in Table 4, we have several findings: (1) *The CDR performance of CDRNP(w/o PRM) significantly outperforms that of CDRNP(w/o PRM+ACP).* The results indicate the effectiveness of transferring the preference of users from the source domain to the target domain, and highlight the necessity of using the ACP module to utilize auxiliary information from the source domain for CDR. (2) *The improvement of CDRNP over CDRNP (w/o PRM) is significant.*


**Figure 4: Performance comparison of CDR to cold-start users with different lengths of interaction items history $\mathcal{V}^s(u_j)$.****Figure 5: Performance comparison of CDR for cold-start users with different lengths of support set \mathcal{C}_i .**

The results demonstrate the importance of capturing the preference correlations among users in different layer-depths.

4.5 Impact of $|\mathcal{V}^s(u_j)|$ and $|\mathcal{C}_i|$

In this section, we modify the length of interaction items history $\mathcal{V}^s(u_j)$ and the number of users in \mathcal{C}_i separately in three CDR scenarios with α value of 20% to test the effectiveness of CDRNP.

4.5.1 Impact of $|\mathcal{V}^s(u_j)|$. In this section, $|\mathcal{V}^s(u_j)|$ is chosen from $\{10, 30, 50, 70, 90\}$. From the results shown in Figure 4, we can draw the following conclusions: (1) *Despite a short interaction items history, CDRNP exhibits exceptional recommendation performance.* Our model shows robust performance with regard to the length of interaction items history in the source domain. (2) *CDRNP consistently outperforms CDRNP(w/o PRM) and PTUPCDR when the interaction items history is reduced.* This indicates that our model is more effective in generating accurate recommendations using limited interaction items history in the source domain. (3) *Our model is less affected by the length of the interaction items history.* It is possibly due to its ability to model preference correlations among the overlapping and cold-start users and user-specific preference from both domains, resulting in complete preferences of users. In contrast, PTUPCDR is more sensitive to the change of $|\mathcal{V}^s(u_j)|$.

Figure 6: Effect of the training loss factor λ .

4.5.2 Impact of $|C_i|$. In this section, we set $|C_i|$ to $\{20, 30, 40, 50, 60\}$. From the results shown in Figure 5, we can draw the following conclusions: (1) *Both CDRNP and CDRNP (w/o PRM) maintain robust performance when the number of users in the support set C_i is reduced.* This indicates that our model is less affected by the reduction of users in C_i . The stability may be attributed to the fact that CDRNP applies a random function to the constructions of ω_i to simulate different function realisations, making the performance of CDRNP more stable. (2) *CDRNP outperforms CDRNP (w/o PRM) in all experimental settings.* The results further emphasize the importance of capturing preference correlations among the overlapping and cold-start users in different layer-depths.

4.6 Hyper-parameter Analysis

In this section, we investigate the sensitivity of CDRNP to the hyperparameter λ in Eq. (18). This experiment is conducted in three CDR scenarios with α value of 20%. Figure 6(a) shows results of CDRNP, and Figure 6(b) for CDRNP(w/o PRM). From these results, we can draw the following conclusions: (1) *We find λ can be set as a larger value for larger interaction scenario, such as the Book-domain \rightarrow Movie-domain.* The reason might be that the larger interaction scenario contains more complex preference correlations among the overlapping and cold-start users, which should be captured using a larger λ . In contrast, for the other two scenarios, a smaller λ can achieve good recommendation performance. (2) *The experimental results of our model are stable, which verifies the robustness of our model to changes in λ .* Even in the worst case of λ , our model still outperforms other baselines in Table 2.

5 RELATED WORKS

Cross-domain recommendation leverages observed interactions from source domain to improve the recommendation performance in the target domain. Following the collaborative filtering (CF) paradigm, a significant amount of CDR methods [9, 26, 31, 35] have been proposed to make effective recommender systems. Although these CF-based methods achieve outstanding performance, they still suffer from the cold-start and data sparsity issues. As our CDRNP aims to solve the previous issue, we only focus on CDR methods which recommend items to cold-start users.

To mitigate the cold-start issue, several CDR methods have been proposed to establish correlations between the source and target domains. Specifically, EMCDCR [20] follows the embedding and mapping paradigm, which learns user representations in each domain with a mapping function. DCDCSR [38] utilizes MF and DNN to generate benchmark factors from source and target domains. SSCDR [10] proposes a semi-supervised mapping function by exploiting

overlapping users and all items. PTUPCDR [40] and TMCDCR [39] follow the MAML [6] framework to learn a meta network to substitute the mapping function. CATN [37] proposes a model based on user review documents to learn the correlation between the source and target domains. CDRIB [2] attempts to capture the domain-shared information via information bottleneck [32]. Compared with them, CDRNP has important design differences that we capture the preference correlations among overlapping and cold-start users and utilize explicit user-item interactions from both domains.

Meta-learning is a learning-to-learn paradigm, aiming to improve the performance of new tasks through training on a distribution of related tasks. Inspired by this, meta-learning methods have been introduced to alleviate the cold-start problem in recommender systems, which include optimization-based methods [3, 13, 19], gradient-based methods [6, 23, 29] and model-based methods [7, 8, 28]. MeLU [13] fit a candidate selection strategy into a standard model-agnostic meta-learning framework to determine distinct items for customized preference estimation. MAML [6] aims at learning model's initial parameters which are capable of adapting to new tasks with only a few gradient steps. The proposed CDRNP falls into the model-based methods, which utilizes NP [8] to capture the preference correlations among users.

Neural Process is a well-known member of stochastic processes family. On the one hand, NP can model the distribution over functions like Gaussian processes. On the other hand, NP is structured as neural networks which can generate predictions in a efficient way. NP are originally designed to the tasks of few-shot image completion [7, 8], image recognition [34]. Since the first model CNP [7], there have been numerous subsequent works to improve NP in various aspects. ANP [11] introduces a self-attention mechanism to improve the model's reconstruction quality. BNP [14] uses bootstrap to learn the stochasticity in NP without assuming a particular form. Recently, a series of methods [4, 16, 17] utilizing NP have emerged in the field of recommender systems. TANP utilizes Neural Process to solve the cold-start issue. However, TANP focuses on single-domain recommendation, which has quite different task setting comparing to cross-domain recommendation.

6 CONCLUSION

In this paper, we propose a novel meta-learning-based method CDRNP for CDR to cold-start users via Neural Process. CDRNP captures the preference correlations among overlapping and cold-start users with an associated stochastic process. Additionally, CDRNP introduces a preference remainder for enhancing the common preference from the overlapping users, and an adaptive conditional decoder to make prediction. Experimental results demonstrate that CDRNP consistently outperforms SOTA methods in three CDR scenarios. In the future, we will further investigate Neural Process for multi-domain recommendation.

ACKNOWLEDGEMENT

This work was funded by the National Key Research and Development Program of China (No.2021YFB3100600), the Strategic Priority Research Program of Chinese Academy of Sciences (No.XDC02040400), the Youth Innovation Promotion Association of CAS (No.2021153), and the National Postdoctoral Program (No.GZC20232968).

REFERENCES

- [1] Homanga Bharadhwaj. 2019. Meta-learning for user cold-start recommendation. In *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [2] Jiangxia Cao, Jiawei Sheng, Xin Cong, Tingwen Liu, and Bin Wang. 2022. Cross-domain recommendation to cold-start users via variational information bottleneck. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2209–2223.
- [3] Manqing Dong, Feng Yuan, Lina Yao, Xiwei Xu, and Liming Zhu. 2020. Mamo: Memory-augmented meta-optimization for cold-start recommendation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 688–697.
- [4] Jing Du, Zesheng Ye, Bin Guo, Zhiwen Yu, and Lina Yao. 2023. IDNP: Interest Dynamics Modeling using Generative Neural Processes for Sequential Recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 481–489.
- [5] Xinyan Fan, Jianxun Lian, Wayne Xin Zhao, Zheng Liu, Chaozhuo Li, and Xing Xie. 2022. Ada-Ranker: A Data Distribution Adaptive Ranking Paradigm for Sequential Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1599–1610.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. PMLR, 1126–1135.
- [7] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. 2018. Conditional neural processes. In *International conference on machine learning*. PMLR, 1704–1713.
- [8] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. 2018. Neural processes. *arXiv preprint arXiv:1807.01622* (2018).
- [9] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [10] SeongKu Kang, Junyoung Hwang, Dongha Lee, and Hwanjo Yu. 2019. Semi-supervised learning for cross-domain recommendation to cold-start users. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1563–1572.
- [11] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. 2019. Attentive neural processes. *arXiv preprint arXiv:1901.05761* (2019).
- [12] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [13] Hyeon Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. 2019. Melu: Meta-learned user preference estimator for cold-start recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1073–1082.
- [14] Juho Lee, Yoonho Lee, Jungtaek Kim, Eunho Yang, Sung Ju Hwang, and Yee Whye Teh. 2020. Bootstrapping neural processes. *Advances in neural information processing systems* 33 (2020), 6606–6615.
- [15] Pan Li and Alexander Tuzhilin. 2020. Ddtcd: Deep dual transfer cross domain recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 331–339.
- [16] Xixun Lin, Jia Wu, Chuan Zhou, Shirui Pan, Yanan Cao, and Bin Wang. 2021. Task-adaptive neural process for user cold-start recommendation. In *Proceedings of the Web Conference 2021*. 1306–1316.
- [17] Huafeng Liu, Liping Jing, Dahai Yu, Mingjie Zhou, and Michael Ng. 2022. Learning Intrinsic and Extrinsic Intentions for Cold-start Recommendation with Neural Stochastic Processes. In *Proceedings of the 30th ACM International Conference on Multimedia*. 491–500.
- [18] Meng Liu, Jianjun Li, Guohui Li, and Peng Pan. 2020. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 885–894.
- [19] Yuanfu Lu, Yuan Fang, and Chuan Shi. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1563–1573.
- [20] Tong Man, Huawei Shen, Xiaolong Jin, and Xueqi Cheng. 2017. Cross-domain recommendation: An embedding and mapping approach. In *IJCAI*, Vol. 17. 2464–2470.
- [21] Andriy Mnih and Karol Gregor. 2014. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*. PMLR, 1791–1799.
- [22] Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *International conference on machine learning*. PMLR, 2554–2563.
- [23] Alex Nichol and John Schulman. 2018. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999* 2, 3 (2018), 4.
- [24] Bernt Oksendal. 2013. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media.
- [25] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [26] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [27] Aghiles Salah, Thanh Binh Tran, and Hady Lauw. 2021. Towards source-aligned variational models for cross-domain recommendation. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 176–186.
- [28] Ruslan Salakhutdinov, Joshua Tenenbaum, and Antonio Torralba. 2012. One-shot learning with a hierarchical nonparametric bayesian model. In *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*. JMLR Workshop and Conference Proceedings, 195–206.
- [29] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*. PMLR, 1842–1850.
- [30] Jiarui Shi and Quanmin Wang. 2019. Cross-domain variational autoencoder for recommender systems. In *2019 IEEE 11th International Conference on Advanced Infocomm Technology (ICAIT)*. IEEE, 67–72.
- [31] Ajit P Singh and Geoffrey J Gordon. 2008. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 650–658.
- [32] Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (ITW)*. IEEE, 1–5.
- [33] Manasi Vartak, Arvind Thiagarajan, Conrado Miranda, Jeshua Bratman, and Hugo Larochelle. 2017. A meta-learning perspective on cold-start recommendations for items. *Advances in neural information processing systems* 30 (2017).
- [34] Jianfeng Wang, Thomas Lukasiewicz, Daniela Massiceti, Xiaolin Hu, Vladimir Pavlovic, and Alexandros Neophytou. 2022. Np-match: When neural processes meet semi-supervised learning. In *International Conference on Machine Learning*. PMLR, 22919–22934.
- [35] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*. 165–174.
- [36] Jun Xiao, Hao Ye, Xiangnan He, Hanwang Zhang, Fei Wu, and Tat-Seng Chua. 2017. Attentional factorization machines: Learning the weight of feature interactions via attention networks. *arXiv preprint arXiv:1708.04617* (2017).
- [37] Cheng Zhao, Chenliang Li, Rong Xiao, Hongbo Deng, and Aixin Sun. 2020. CATN: Cross-domain recommendation for cold-start users via aspect transfer network. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 229–238.
- [38] Feng Zhu, Yan Wang, Chaochao Chen, Guanfang Liu, Mehmet Orgun, and Jia Wu. 2020. A deep framework for cross-domain and cross-system recommendations. *arXiv preprint arXiv:2009.06215* (2020).
- [39] Yongchun Zhu, Kaikai Ge, Fuzhen Zhuang, Ruobing Xie, Dongbo Xi, Xu Zhang, Leyu Lin, and Qing He. 2021. Transfer-meta framework for cross-domain recommendation to cold-start users. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1813–1817.
- [40] Yongchun Zhu, Zhenwei Tang, Yudan Liu, Fuzhen Zhuang, Ruobing Xie, Xu Zhang, Leyu Lin, and Qing He. 2022. Personalized transfer of user preferences for cross-domain recommendation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 1507–1515.