Contrastive Cross-domain Recommendation in Matching

Ruobing Xie*
WeChat, Tencent
Beijing, China
ruobingxie@tencent.com

Qi Liu*
WeChat, Tencent
Beijing, China
addisliu@tencent.com

Liangdong Wang
WeChat, Tencent
Beijing, China
ldwang@tencent.com

Shukai Liu
WeChat, Tencent
Beijing, China
shukailiu@tencent.com

Bo Zhang
WeChat, Tencent
Beijing, China
nevinzhang@tencent.com

Leyu Lin
WeChat, Tencent
Beijing, China
goshawklin@tencent.com

arXiv:2112.00999v2 [cs.IR] 21 Jun 2022

ABSTRACT

Cross-domain recommendation (CDR) aims to provide better recommendation results in the target domain with the help of the source domain, which is widely used and explored in real-world systems. However, CDR in the matching (i.e., candidate generation) module struggles with the data sparsity and popularity bias issues in both representation learning and knowledge transfer. In this work, we propose a novel Contrastive Cross-Domain Recommendation (CCDR) framework for CDR in matching. Specifically, we build a huge diversified preference network to capture multiple information reflecting user diverse interests, and design an intra-domain contrastive learning (intra-CL) and three inter-domain contrastive learning (inter-CL) tasks for better representation learning and knowledge transfer. The intra-CL enables more effective and balanced training inside the target domain via a graph augmentation, while the inter-CL builds different types of cross-domain interactions from user, taxonomy, and neighbor aspects. In experiments, CCDR achieves significant improvements on both offline and online evaluations in a real-world system. Currently, we have deployed our CCDR on WeChat Top Stories, affecting plenty of users. The source code is in https://github.com/lqfarmer/CCDR.

CCS CONCEPTS

• **Information systems → Recommender systems**.

KEYWORDS

contrastive learning, cross-domain recommendation, matching

ACM Reference Format:
Ruobing Xie, Qi Liu, Liangdong Wang, Shukai Liu, Bo Zhang, and Leyu Lin. 2022. Contrastive Cross-domain Recommendation in Matching. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22), August 14–18, 2022, Washington, DC, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3534678.3539125

*Both authors have equal contributions. Ruobing Xie is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD '22, August 14–18, 2022, Washington, DC, USA
© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9385-0/22/08…$15.00
https://doi.org/10.1145/3534678.3539125

1 INTRODUCTION

Personalized recommendation aims to provide attractive items for users according to their profiles and historical behaviors, which has been widely implemented in various fields of our lives. Real-world large-scale recommendation systems usually adopt the classical two-stage architecture containing ranking and matching. The *matching* module [35, 36, 38] (i.e., candidate generation [3]) focuses more on the efficiency and diversity, which first retrieves a small subset of (usually hundreds of) item candidates from the million-level large corpora. Next, the *ranking* module gives the specific ranks of items for the final display.

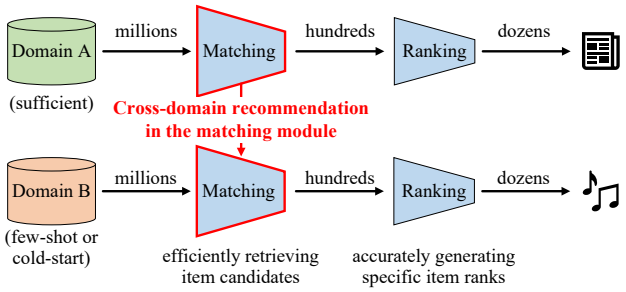

Figure 1: An example of CDR in matching.

With the increase of recommendation scale and the expansion of recommendation scenarios, real-world recommendations usually need to bring in additional data sources (i.e., domains) as supplements to improve their content coverage and diversity. These cold-start items of new data sources only have very few user behaviors at their warm-up stage. Hence, it is difficult to recommend these cold-start items appropriately. *Cross-domain recommendation (CDR)*, which aims to make full use of the informative knowledge from the source domain to help the target domain's recommendation [47], is proposed to solve this issue. EMCDR [17] is a classical CDR method, which focuses on building user mapping functions via aligned user representations in the source and target domains. CoNet [9] proposes another approach that jointly models feature interactions in two domains via a cross connection unit. However, existing CDR methods often heavily rely on aligned users for cross-domain mapping (e.g., EMCDR), ignoring other rich information in recommendation such as taxonomy. It will harm the knowledge transfer between different domains, especially in cold-start scenarios. Moreover, lots of CDR methods are designed for ranking

that consider complicated cross-domain user-item interactions (e.g., CoNet), which cannot be directly adopted in matching due to the online efficiency. CDR in the matching module should consider not only recommendation accuracy, but also diversity and efficiency.

In this work, we aim to improve the matching module's performance on new (few-shot or strict cold-start) domains via the CDR manner. Fig. 1 shows an illustration of this task. Precisely, CDR in matching mainly has the following three challenges:

(1) *How to address the data sparsity and popularity bias issues of CDR in matching?* Real-world recommendation usually suffers from serious data sparsity issues when modeling the interactions between million-level users and items. Moreover, these sparse interactions are even highly skewed to popular items with high exposure owing to the Matthew effect [20], which makes hot items become hotter. These two issues inevitably harm the representation learning of cold-start and long-tail items, whose damages will even be multiplied in matching where all items should be considered.

(2) *How to conduct more effective knowledge transfer for the (cold-start) target domain with few user behaviors?* As stated above, conventional CDR methods strongly depend on aligned users and their behaviors. The performance of CDR in matching will be greatly reduced, if most users and items have few interactions and models cannot learn reliable representations in cold-start domains. Moreover, other heterogeneous information (e.g., taxonomy) should also be fully considered in CDR to bridge different domains. We should build more effective and robust cross-domain knowledge transfer paths to well learn both popular and long-tail objects.

(3) *How to balance the practical demands of accuracy, diversity and efficiency of CDR in matching?* Online efficiency requirements need to be strictly followed. Moreover, matching is more responsible for the diversity than ranking, for it determines the inputs of ranking. A good CDR matching model should comprehensively transfer user diverse preferences via multiple paths to the target domain.

To address these issues, we propose a novel **Contrastive Cross-Domain Recommendation (CCDR)** to transfer user preferences in matching. Specifically, we build two global diversified preference networks for two domains, containing six types of objects to enhance diversity and cross-domain connections. We conduct a GNN aggregator with a neighbor-similarity based loss on heterogeneous interactions to capture user diverse interests. To strengthen the cross-domain knowledge transfer, we design the **intra-domain contrastive learning (intra-CL)** and **inter-domain contrastive learning (inter-CL)** in CCDR. The intra-CL conducts an additional self-supervised learning with sub-graph based data augmentations to learn more reliable representations for matching in the target domain. In contrast, the inter-CL designs three contrastive learning tasks focusing on the cross-domain mapping between aligned users, taxonomies, and their heterogeneous neighbors. The mutual information maximization with different types of objects multiplies the effectiveness of cross-domain knowledge transfer. Finally, all additional CL losses are combined with the original CDR losses under a multi-task learning (MTL) framework. We conduct a cross-domain multi-channel matching to further improve the diversity in online. CCDR has the following three advantages: (1) The intra-CL brings in self-supervised learning for long-tail users and items, which can alleviate the data sparsity and popularity bias issues in matching. (2) The inter-CL introduces new CL-based cross-domain interactions, which enables more effective and robust knowledge transfer for CDR in matching. (3) The diversified preference network, multiple CL tasks, and the cross-domain multi-channel matching cooperate well to capture user diverse preferences, which meets the requirements of diversity and efficiency in online system.

In experiments, we compare CCDR with competitive baselines on real-world recommendation domains, and achieve significant improvements on both offline and online evaluations. Moreover, we also conduct some ablation tests and parameter analyses to better understand our model. The contributions are concluded as:

- We propose a novel contrastive cross-domain recommendation for CDR in matching. To the best of our knowledge, we are the first to conduct contrastive learning to improve representation learning and knowledge transfer in CDR.
- We propose the intra-CL task with a sub-graph based data augmentation to learn better node representations inside the single domain, which can alleviate the data sparsity issue in CDR of matching.
- We also creatively design three inter-CL tasks via aligned users, taxonomies, and their neighbors in our diversified preference networks, which enable more effective and diverse knowledge transfer paths across different domains.
- We achieve significant improvements on both offline and online evaluations. CCDR is simple, effective, and easy to deploy. It has been deployed on a real-world recommendation system for more than 6 months, affecting plenty of users.

2 RELATED WORKS

Matching in Recommendation. The matching (i.e., candidate generation) module aims to retrieve a small subset of (usually hundreds of) items from large corpora efficiently [36]. Recently, embedding-based retrieval [1, 10] is also widely used in practical systems with the help of fast retrieval servers [11]. Due to the need for efficiency, embedding-based methods usually adopt a two-tower architecture, which conducts two different towers for building user and item representations separately. Different feature interaction modeling methods such as FM [23], YoutubeDNN [3], AutoInt [25], ICAN [36], and AFT [5] could be used in these towers. In contrast, tree-based matching models [46] give another way to address the matching problem with structured item trees. Graph-based matching models [35] are also proposed to learn user/item representations. However, few works focus on the CDR in matching, which often exists in practical recommendations.

Cross-domain Recommendation. Cross-domain recommendation attempts to learn useful knowledge from the source domain to help the target domain's recommendation. EMCDR [17] is a classical embedding mapping approach, which builds the mapping function via aligned users' representations. SSCDR [12] designs a semi-supervised manner to learn item mapping based on EMCDR. In contrast, CoNet [9] is another classical type of CDR method that uses the cross connection unit to model domain interactions. Zhao et al. [42] combines items of both source and target domains in one graph to learn representations. Dual transfer [14], source-target mixed attention [19], and meta networks [47, 48] are also proposed for CDR. Additional review information is also used to enhance the attentive knowledge transfer [43]. However, most of these CDR

models are specially designed for the ranking module (which involve user-item interactions in cross-domain modeling). ICAN [36] is the most related work, which captures field-level feature interactions to improve matching in multiple domains. In CCDR, we introduce several novel CL tasks. To the best of our knowledge, we are the first to conduct CL to jointly improve representation learning and knowledge transfer in CDR.

Contrastive Learning. Contrastive learning (CL) is a representative self-supervised learning (SSL) method, which aims to learn models by contrasting positive pairs against negative pairs. MoCo [6] builds a large dynamic dictionary with a queue and a moving-averaged momentum encoder. SimCLR [2] designs a simple contrastive learning framework with a composition of data augmentations and projectors for CL. BYOL [4] relies on its online and target networks, which iteratively bootstraps the outputs of a network to serve as targets for learning. Some works also consider graph contrastive learning [22, 39].

CL in recommendation. Recently, SSL and CL are also verified in recommendation [28, 41]. S^3-Rec [45] builds contrastive learning tasks among items, attributes, sentences, and sub-sentences in sequential recommendation. UPRec focuses on user-aware SSL [34]. Wu et al. [31] adopts CL between behaviors and models in multi-behavior recommendation. Moreover, CL has also been used in disentangled recommendation [44], social recommendation [40], sequential recommendation [32, 33, 37]. For graph-based CL, Wu et al. [30] introduces embedding, node, edge dropouts to graph-based recommendation. Differing from these works, we build three CL tasks to facilitate the user preference transfers between different domains in cold-start cross-domain recommendation.

3 METHODOLOGY

In this work, we propose CCDR to enhance the cross-domain recommendation in matching via contrastive learning.

3.1 Problem Definition and Overall Framework

CDR in matching. We concentrate on the matching module of the classical two-stage recommendation systems [3]. Matching is the first step before ranking, which attempts to efficiently retrieve hundreds of items from million-level item candidates. It cares more about whether good items are retrieved (often measured by hit rate), not the specific top item ranks which should be considered by the following ranking module (often measured by NDCG or AUC) [16, 36]. The CDR in matching task attempts to improve the target domain's matching module with the help of the source domain.

Overall framework. CCDR is trained with three types of losses, including the original source/target single-domain losses, the intra-domain CL loss, and the inter-domain CL loss. (1) We first build a huge global diversified preference network separately for each domain as the sources of user preferences. This diversified preference network contains various objects such as user, item, tag, category, media, and word with their interactions to bring in user diverse preferences from different aspects. (2) Next, we train the single-domain matching model via a GNN aggregator and the neighbor-similarity based loss. (3) Since the cold-start domain lacks sufficient user behaviors, we introduce the intra-domain CL inside the target domain to train more reliable node representations with a sub-graph based data augmentation. (4) To enhance the cross-domain knowledge transfer, we design three inter-domain CL tasks via aligned users, taxonomies, and their neighbors between two domains, which cooperate well with the diversified preference network. All three losses are combined under a multi-task learning framework.

3.2 Diversified Preference Network

Conventional matching [3] and CDR [12, 17] models usually heavily rely on user-item interactions to learn CTR objectives and cross-domain mapping. However, it will decrease the diversity of matching due to the popularity bias issue. Moreover, it does not take full advantage of other connections (e.g., tags, words, medias) besides users between different domains, which is particularly informative in cross-domain knowledge transfer.

Therefore, inspired by [15, 35], we build a global diversified preference network for each domain, considering 6 types of important objects in recommendation as nodes and their heterogeneous interactions as edges. Specifically, we use *item, user, tag, category, media, and word* as nodes. Tags and categories are item taxonomies that represent users' fine- and coarse- granularity interests. Media indicates the item's producer. Words reflect the semantic information of items extracted from items' titles or contents. To alleviate data sparsity and accelerate our offline training, we also gather users into user groups according to their basic profiles (all users having the same gender-age-location attributes are clustered in the same user group). These user groups are viewed as user nodes in CCDR.

As for edges, we consider the following six item-related interactions: (a) *User-item edge (U-I).* This edge is generated if an item is interacted by a user group at least 3 times. We jointly consider multiple user behaviors (i.e., click, like, share) to build this edge with different weights. (b) *Item-item edge (I-I).* The I-I edge introduces sequential information of user behaviors in sessions. It is built if two items appear in adjacent positions in a session. (c) *Tag-item edge (T-I).* The T-I edges connect items and their tags. It captures items' fine-grained taxonomy information. (d) *Category-item edge (C-I).* It records items' coarse-grained taxonomy information. (e) *Media-item edge (M-I).* It links items with their producers/sources. (f) *Word-item edge (W-I).* It highlights the semantic information of items from their titles. Each edge is undirected but empirically weighted according to the edge type and strength (e.g., counts for U-I edges). Compared with conventional U-I graphs, our diversified preference network tries its best to describe items from different aspects via these heterogeneous interactions. The advantages of this diversified preference network are: (1) it brings in additional information as supplements to user-item interactions, which jointly improve accuracy and diversity (Sec. 3.3.2). (2) It can build more potential bridges between different domains via users, tags, categories, and words, which cooperates well with the inter-CL tasks and the online multi-channel matching in CDR (Sec. 3.5 and Sec. 4).

3.3 Single-domain GNN Aggregator

3.3.1 GNN-based Aggregator. Inspired by the great successes of GNN, we adopt GAT [27] as the aggregator on the diversified preference network for simplicity and universality. Precisely, we randomly initialize $\boldsymbol{e}_i^0$ for all heterogeneous nodes. For a node e_i and its neighbor $e_k \in N_{e_i}$ (N_{e_i} is the neighbor set of e_i after a weighted

sampling), we have e_i's node representation $\boldsymbol{e}_i^x$ at the x-th layer as:

$$\boldsymbol{e}_i^x = \sigma(\sum_{e_k \in N_{e_i}} \alpha_{ik}^x \boldsymbol{W}^x \boldsymbol{e}_k^{x-1}). \tag{1}$$

$\boldsymbol{W}^x$ is the weighting matrix, σ is the sigmoid function. α_{ik}^x represents the attention between e_i and e_k in x-th layer noted as:

$$\alpha_{ik}^x = \frac{\exp(f(\boldsymbol{w}^{x^\top}[\boldsymbol{W}^x \boldsymbol{e}_i^{x-1} || \boldsymbol{W}^x \boldsymbol{e}_k^{x-1}]))}{\sum_{e_l \in N_{e_i}} \exp(f(\boldsymbol{w}^{x^\top}[\boldsymbol{W}^x \boldsymbol{e}_i^{x-1} || \boldsymbol{W}^x \boldsymbol{e}_l^{x-1}]))}, \tag{2}$$

where $f(\cdot)$ indicates a LeakyReLU activation and $||$ indicates the concatenation. $\boldsymbol{w}^x$ is the weighting vector. Note that the N_{e_i} is a dynamic neighbor set which is randomly generated based on the edge weight in Sec. 3.2. We conduct a two-layer GAT to generate the aggregated node representations $\boldsymbol{e}_i$ for all nodes ($\boldsymbol{e}_i^s$ and $\boldsymbol{e}_i^t$ for the source and target domains). It is also not difficult to conduct other GNN models such as LightGCN [7] in this module.

3.3.2 Neighbor-similarity Based Optimization. In practical CDR scenarios, users often have fewer historical behaviors on items in (cold-start) target domains. Conventional embedding-based matching methods such like Matrix factorization (MF) [13] cannot get sufficient supervised information from the sparse user-item interactions, and thus cannot learn reliable user and item representations for matching. To capture additional information from behavior, session, taxonomy, semantics, and data source aspects, we conduct the neighbor-similarity based loss [15] on the diversified preference network. As shown in Fig. 2, this loss projects all nodes into the same latent space, making all nodes similar with their neighbors. It regards all types of edges as unsupervised information to guide the training besides user-item interactions. Formally, the neighbor-similarity based loss L_N is defined as follows:

$$L_N = -\sum_{e_i} \sum_{e_k \in N_{e_i}} \sum_{e_j \notin N_{e_i}} (-\log(\sigma(\boldsymbol{e}_i^\top \boldsymbol{e}_j)) + \log(\sigma(\boldsymbol{e}_i^\top \boldsymbol{e}_k))). \tag{3}$$

$\boldsymbol{e}_i$ is the i-th aggregated node representation, and e_k is a sampled neighbor of e_i. e_j is a randomly selected negative sample of e_i.

We choose the neighbor-similarity based loss for the following advantages: (1) L_N makes full use of all types of interactions between heterogeneous objects in matching, which contain significant information from user behaviors (U-I edges), sessions (I-I edges), item taxonomies (T-I and C-I edges), data sources (M-I edges) and semantics (W-I edges). It helps to capture user diverse preferences to balance accuracy and diversity in matching. If we only consider U-I edges, this loss will degrade into the classical MF. (2) CDR in matching should deal with long-tail items. L_N brings in additional information for long-tail items that can benefit cold-start domains. (3) We conduct a cross-domain multi-channel matching strategy in online for diversity. This embedding-based retrieval strategy also depends on heterogeneous node embeddings optimized by L_N to retrieve similar items in the (cold-start) target domain (see Sec. 4 for more details). The L_N loss exactly fits the online multi-channel matching, and also well cooperates with the diversified preference network and the inter-CL losses. We cannot conduct complicated user-item interaction calculations in Eq. (3), since we rely on the fast embedding-based retrieval in matching for efficiency.

3.4 Intra-domain Contrastive Learning

Contrastive learning is a widely-used SSL method that can make full use of unlabelled data via its pair-wise training. In CCDR, we conduct two types of CL tasks. The **intra-domain contrastive learning (intra-CL)** is conducted inside the target domain to learn better node representations, while the **inter-domain contrastive learning (inter-CL)** is adopted across the source and target domains to guide a better knowledge transfer.

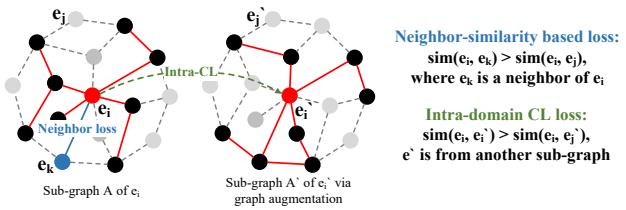

Figure 2: The neighbor-similarity loss and the intra-CL loss.

In intra-CL, we conduct a sub-graph based data augmentation for each node aggregation, which could be regarded as a dynamic node/edge dropout in classical graph augmentation [30]. Precisely, for a node e_i, we sample two neighbor set N_{e_i} and N'_{e_i} to conduct the GNN aggregation, and receive two node representations $\boldsymbol{e}_i$ and $\boldsymbol{e}'_i$. $\boldsymbol{e}'_i$ is regarded as the positive instance of $\boldsymbol{e}_i$ in intra-CL, with a different sub-graph sampling focusing on different neighbors of e_i. Similar to [2], we randomly sample from other examples $\boldsymbol{e}'_j$ in the same batch B of e_i to get the negative samples e_j. We do not use all examples in B as negative samples for efficiency. In this case, the popularity bias is partially solved [29]. Formally, we follow the InfoNCE [18] to define the intra-CL loss L_{intra} as follows:

$$L_{intra} = -\sum_B \sum_{e_i \in B} \log \frac{\exp(\text{sim}(\boldsymbol{e}_i, \boldsymbol{e}'_i)/\tau)}{\sum_{e'_j \in S_{Bi}} \exp(\text{sim}(\boldsymbol{e}_i, \boldsymbol{e}'_j)/\tau)}. \tag{4}$$

S_{Bi} indicates the negative samples of e_i in B. τ is the temperature. $\text{sim}(\boldsymbol{e}_i, \boldsymbol{e}'_j)$ measures the similarity between $\boldsymbol{e}_i$ and $\boldsymbol{e}'_j$, which is calculated as their cosine similarity. With the intra-CL loss, long-tail nodes can also get training opportunities via SSL.

3.5 Inter-domain Contrastive Learning

The inter-CL aims to improve the knowledge transfer across different domains via various types of nodes and edges in the diversified preference network. Precisely, we design three inter-domain CL tasks via aligned users, taxonomies, and neighbors as in Fig. 3.

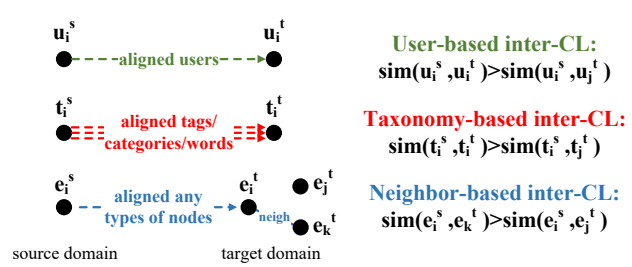

Figure 3: Three inter-CL tasks across different domains.

3.5.1 User-based Inter-CL. Most conventional CDR methods [17] take aligned users as their dominating mapping seeds across domains. We follow this idea and conduct a user-based inter-CL task. Each user u_i has two user representations $\boldsymbol{u}_i^s$ and $\boldsymbol{u}_i^t$ in the source and target domains learned in Sec. 3.3. Although users may have different preferences and behavior patterns in two domains, it is still natural that the source-domain representation $\boldsymbol{u}_i^s$ should be more similar with its target-domain $\boldsymbol{u}_i^t$ than any other representations $\boldsymbol{u}_j^t$. We define the user-based inter-CL loss L_{inter_u} as follows:

$$L_{inter_u} = -\sum_{u_i} \log \frac{\exp(\text{sim}(\boldsymbol{u}_i^s, \boldsymbol{u}_i^t)/\tau)}{\sum_{u_j \in S_{u_i}} \exp(\text{sim}(\boldsymbol{u}_i^s, \boldsymbol{u}_j^t)/\tau)}. \tag{5}$$

S_{u_i} is the sampled negative set collecting from all other users except u_i. $\text{sim}(\cdot,\cdot)$ indicates the cosine similarity.

3.5.2 Taxonomy-based inter-CL. Differing from some classical CDR methods [12], CCDR builds a diversified preference network that introduces more bridges across different domains. We assume that the same tag/category/word in different domains should have the same meanings. Hence, we design a taxonomy-based inter-CL similar to the user-based CL. we take the aggregated node representation pair $(\boldsymbol{t}_i^s, \boldsymbol{t}_i^t)$ of the same taxonomy t_i in two domains as the positive pair, where t_i could be tags, categories, and words. We have:

$$L_{inter_t} = -\sum_{t_i} \log \frac{\exp(\text{sim}(\boldsymbol{t}_i^s, \boldsymbol{t}_i^t)/\tau)}{\sum_{t_j \in S_{t_i}} \exp(\text{sim}(\boldsymbol{t}_i^s, \boldsymbol{t}_j^t)/\tau)}, \tag{6}$$

S_{t_i} is the sampled negative set of t_i from all other taxonomies with the same type. We can set different temperatures for taxonomies and users if we want to sharpen the differences of some types. L_{inter_t} functions as a supplement to the original user-based mapping.

3.5.3 Neighbor-based inter-CL. Besides the explicit alignments of users and taxonomies across domains, there are also some essential objects such as items that do not have explicit mapping. We aim to bring in more implicit cross-domain knowledge transfer paths between unaligned nodes in two domains. We suppose that similar nodes in different domains should have similar neighbors (e.g., similar items may have similar users, taxonomies, and producers). Hence, we propose a neighbor-based inter-CL, which builds indirect (multi-hop) connections between objects in different domains. Precisely, we define E_A as the overall aligned node set (including users, tags, categories, and words). The neighbor-based inter-CL loss L_{inter_n} is formalized with all aligned nodes $e_i \in E_A$ and e_i's neighbor set $N_{e_i}^t$ in the target domain as follows:

$$L_{inter_n} = -\sum_{e_i \in E_A} \sum_{e_k \in N_{e_i}^t} \log \frac{\exp(\text{sim}(\boldsymbol{e}_i^s, \boldsymbol{e}_k^t)/\tau)}{\sum_{e_j \notin N_{e_i}^t} \exp(\text{sim}(\boldsymbol{e}_i^s, \boldsymbol{e}_j^t)/\tau)}. \tag{7}$$

In L_{inter_n}, for an aligned node's representation $\boldsymbol{e}_i^s$ in the source domain, its target-domain neighbor's representation $\boldsymbol{e}_k^t$ is the positive instance, while other target-domain representations $\boldsymbol{e}_j^t$ are negative. It is reasonable since related objects should be connected in the diversified preference network and learned to be similar under the neighbor-similarity based loss in Eq. (3). It is also convenient to extend the current positive samples $e_k \in N_{e_i}^t$ to multi-hop neighbors for better generalization and diversity in CDR.

This neighbor-based inter-CL greatly multiplies the diversified knowledge transfer paths between two domains, especially for the cold-start items. For example, through the $tag_i^s \rightarrow tag_i^t \rightarrow item_j^t$ path, the cold-start $item_j$'s representation in the target domain can be directly influenced by fully-trained representations in the source domain. Moreover, the similarities between different types of source-domain node representations and the target-domain item representations are directly used in the online multi-channel matching for diversified retrieval, which will be discussed in Sec. 4.

Finally, we combine all three CL losses from aligned user, taxonomy, and neighbor aspects to form the inter-CL loss L_{inter} as:

$$L_{inter} = L_{inter_u} + L_{inter_t} + L_{inter_n}. \tag{8}$$

3.6 Multi-task Optimization

Following classical CL-based recommendation models [40], we also conduct a multi-task optimization combining the source-domain matching loss L_{N_s}, the target-domain matching loss L_{N_t}, the intra-CL loss L_{intra}, and the inter-CL loss L_{inter} as follows:

$$L = \lambda_1 L_{N_s} + \lambda_2 L_{N_t} + \lambda_3 L_{intra} + \lambda_4 L_{inter}. \tag{9}$$

$\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are loss weights set as 1.0, 1.0, 1.5, 0.6 according to the grid search. More details and motivations are in Appendix A.3.

4 ONLINE DEPLOYMENT

We have deployed CCDR on the cold-start matching module in a well-known recommendation system named WeChat Top Stories. A good CDR-based cold-start matching module should have the following key characteristics: (1) making full use of user behaviors and item features in the source and target domains, (2) capturing user diverse preferences from different aspects, and (3) balancing accuracy, diversity and efficiency. To achieve these, we propose a new **cross-domain multi-channel matching** in online.

Specifically, we conduct six channels including *user*, *item*, *tag*, *category*, *media*, and *word* channels to retrieve items in the target domains via node representations learned by Eq. (9). We rely on the user historical behavior sequence $seq = \{d_1, d_2, \cdots, d_n\}$ to capture user's interests, where d_i is the i-th clicked item and n is the max length. In the item channel, we directly use the node representations of all items in seq to retrieve similar items in the target domain. Formally, we define the score s_i^d of the i-th target-domain item $\bar{d}_i$ in the item channel as follows:

$$s_i^d = \sum_{j=1} sim(\bar{\boldsymbol{d}}_i, \boldsymbol{d}_j) \times satisf_j \times recency_j \times z^d(i,j). \tag{10}$$

$sim(\bar{\boldsymbol{d}}_i, \boldsymbol{d}_j)$ is the cosine similarity between the clicked item d_j in user historical behaviors and the item candidate $\bar{d}_i$ in the target domain, where $\bar{\boldsymbol{d}}_i$ and $\boldsymbol{d}_j$ are aggregated item embeddings trained by Eq. (9). $satisf_j$ measures the posterior user satisfaction on d_j, which is calculated as user's complete rate on the item. $recency_j$ models the temporal factors of historical items, which decays exponentially from the short term to the long term ($recency_j = 0.95^{n-j}$). For online efficiency, each item in seq only recommends its top 100 nearest items. $z^d(i,j)$ equals 1 only if the target-domain item $\bar{d}_i$ appears in the top 100 nearest items of d_j, and otherwise $z^d(i,j) = 0$. We pre-calculate the similarities and index the top nearest items for all nodes in offline to further accelerate the online matching.

To capture user diverse preferences from different aspects, we further conduct the tag, category, media and word channels similar to the item channel. Taking the tag channel for instance, we build a historical tag sequence $seq_t = \{T_1, T_2, \cdots, T_n\}$ according to the item sequence $\{d_1, d_2, \cdots, d_n\}$, where T_j is the tag set of d_j. All tags in seq_t retrieve top 100 nearest items in the target domains as candidates. Similar to Eq. (10), the score of the i-th target-domain item $\bar{d}_i$ in the tag channel s_i^t is defined as follows:

$$s_i^t = \sum_{j=1}^{n} \sum_{t_k} sim(\bar{\boldsymbol{d}}_i, \boldsymbol{t}_k) \times satisf_j \times recency_j \times z^t(i, j, k). \quad (11)$$

$sim(\bar{\boldsymbol{d}}_i, \boldsymbol{t}_k)$ is the cosine similarity between $\bar{\boldsymbol{d}}_i$ and the aggregated tag representation $\boldsymbol{t}_k$. $z^t(i, j, k)$ is the tag's indicator. $z^t(i, j, k) = 1$ only if the tag t_k belongs to the j-th item d_j in seq, and $\bar{d}_i$ locates in the top 100 nearest items of t_k. Other category, media and word channels are the same as the tag channel, generating their corresponding scores s_i^c, s_i^m and s_i^w. As for the user channel, we directly depend on the user's gender-age-location attribute triplet's (i.e., the user group in Sec. 3.2) node representations to retrieve top nearest items according to the cosine similarity score s_i^u for $\bar{d}_i$.

Finally, all top items retrieved by six heterogeneous channels are combined and reranked via the aggregated score s_i as:

$$s_i = s_i^u + s_i^d + s_i^t + s_i^c + s_i^m + s_i^w. \quad (12)$$

It is easy to set and adjust the hyper-parameters of heterogeneous channels' weights for the practical demands and the preferences of systems. We rank top target-domain items via s_i, and select top 500 items as the final output of our multi-channel matching, considering both matching accuracy and memory/computation costs.

We conclude the feasibility and advantages of our cross-domain multi-channel matching as follows: (1) These multiple matching channels rely on the similarities between the target-domain items and heterogeneous nodes, which is consistent with the neighbor-similarity based loss and the inter-CL losses. (2) The multi-channel matching makes full use of all heterogeneous information to generate diversified item candidates, which is essential in cold-start matching. (3) We pre-calculate the indexes for the top nearest items of all nodes, which greatly reduces the online computation costs. The online computation complexity of CCDR is $O(\log(600n)+600n)$ (n is the length of user historical behavior). More details of online deployment and efficiency are given in Appendix A.1.

5 EXPERIMENTS

5.1 Large-scale CDR Matching Dataset

CCDR relies on item-related taxonomy, semantic, and producer information for CDR in matching, while no large-scale public CDR dataset is capable for this setting. Therefore, we build a new CDR dataset CDR-427M extracted from a real-world recommendation system named WeChat Top Stories, which contains a source domain and two target domains. Specifically, we randomly select nearly 63 million users, and collect their 427 million behaviors on 3.0 million items. We split these behaviors into the train set and the test set using the chronological order. We also bring in 187 thousand tags, 356 categories, 56 thousand medias, and 207 thousand words as additional item-related information. All data are preprocessed via data masking to protect the user's privacy.

To simulate different CDR scenarios, we evaluate on two target domains having different cold-start degrees. The first is a few-shot target domain, where most users only have several behaviors. The second is a strict cold-start domain, which is more challenging since all user behaviors on items in the train set are discarded [21]. Following Sec. 3.2, we build three diversified preference networks for all domains separately via the train set and all item-related information. More details are in Table 1 and Appendix A.2.

Table 1: Statistics of three domains in CDR-427M.

domain	user	item	edge	train	test
source domain	63.1M	2.04M	175M	406M	/
few-shot target domain	2.38M	0.39M	29.8M	10.8M	4.99M
strict cold-start domain	2.23M	0.57M	15.0M	/	5.48M

5.2 Competitors

We implement several competitive baselines focusing on the matching module and cross-domain recommendation for comparisons.

Classical Matching Methods. We implement three competitive matching models as baselines. We do not compare with tree-based models [46], for they cannot be deployed in cold-start domains. All user behaviors of two domains are considered in these models.

- **Factorization Machine (FM) [Rendle 2010].** FM [23] is a classical embedding-based matching model. It captures the feature interactions between users and items for embedding-based retrieval under the two-tower architecture [3].
- **AutoInt [Song et al. 2019].** AutoInt [25] is a recent method that utilizes self-attention to model feature interactions. It also adopt the two-tower architecture for matching.
- **GraphDR+ [Xie et al. 2021a].** GraphDR [35] is an effective graph-based matching model. The single-domain model of CCDR could be considered as an enhanced GraphDR with differences in node aggregation and multi-channel matching specially designed for CDR. We directly conduct the single-domain model of CCDR on the joint network containing both source and target domains, noted as GraphDR+.

Cross-domain/Multi-domain Methods. We also implement two representative CDR models and one multi-domain matching model as baselines. We do not compare with CDR models like CoNet [9], since they cannot be directly used in matching for efficiency.

- **EMCDR+ [Man et al. 2017].** EMCDR [17] is a classical CDR model that directly learns the embedding mapping of users between two domains. For fair comparisons, we use the same single-domain model and multi-channel matching in CCDR for learning and serving, noted as EMCDR+.
- **SSCDR+ [Kang et al. 2019].** SSCDR [12] is regarded as an enhanced EMCDR, which adopts a semi-supervised loss to learn the mapping of items. We also follow the same settings of EMCDR to get SSCDR+. Since the strict cold-start domain has no user-item behaviors, we use aligned taxonomies to learn cross-domain mappings in EMCDR+ and SSCDR+.

Table 2: Results of matching-related metrics on CDR-427M. All improvements of CCDR are significant (t-test with $p < 0.01$). * indicates that these models are based on the same single-domain model (noted as GraphDR+) in CCDR.

Model	few-shot target domain				strict cold-start target domain			
	HIT@50	HIT@100	HIT@200	HIT@500	HIT@50	HIT@100	HIT@200	HIT@500
FM (Rendle 2010)	0.0101	0.0176	0.0293	0.0551	0.0048	0.0084	0.0145	0.0299
AutoInt (Song et al. 2019)	0.0125	0.0219	0.0370	0.0705	0.0058	0.0102	0.0173	0.0356
GraphDR+ (Xie et al. 2021a)*	0.0132	0.0236	0.0402	0.0780	0.0087	0.0156	0.0266	0.0543
EMCDR+ (Man et al. 2017)*	0.0137	0.0242	0.0414	0.0801	0.0091	0.0160	0.0274	0.0561
SSCDR+ (Kang et al. 2019)*	0.0146	0.0261	0.0448	0.0879	0.0104	0.0183	0.0312	0.0637
ICAN (Xie et al. 2020)	0.0163	0.0291	0.0502	0.0994	0.0074	0.0131	0.0226	0.0464
Sub-graph CL (GraphDR+)*	0.0174	0.0312	0.0541	0.1099	0.0118	0.0209	0.0361	0.0732
Cross-domain KD (GraphDR+)*	0.0169	0.0305	0.0530	0.1058	0.0131	0.0229	0.0387	0.0780
CCDR (inter-CL)*	0.0213	0.0378	0.0671	0.1327	0.0154	0.0271	0.0467	0.0931
CCDR (inter-CL+intra-CL)*	**0.0225**	**0.0411**	**0.0715**	**0.1442**	**0.0161**	**0.0284**	**0.0487**	**0.0987**

- **ICAN [Xie et al. 2020].** ICAN [36] is the SOTA model in multi-domain matching, which is the most related work of our task. It highlights the interactions between feature fields in different domains for cold-start matching.

Knowledge Distillation/Contrastive Learning Methods. We further propose two enhanced versions of the single-domain matching model in CCDR (i.e., GraphDR+) for more challenging comparisons.

- **Sub-graph CL.** We build a sub-graph CL method based on GraphDR+. It considers the intra-CL loss with a sub-graph augmentation in Eq. (4) inspired by [22, 30]. It can be viewed as an ablation version of CCDR without the inter-CL.
- **Cross-domain KD.** We further propose a cross-domain knowledge distillation (KD) model. This model also follows the single-domain model of CCDR, learning the cross-domain mapping via the Hint loss [24] between aligned nodes in two domains (i.e., user, tag, category, and word).

5.3 Experimental Settings

In the single-domain model of CCDR, the input dimensions of all nodes are 128, and the output dimensions are 100. We conduct a weighted neighbor sampling to select 25 and 10 neighbors for the first and second layers' aggregations. The edge weight is proportional to the mutual information between its two nodes to make sure that different types of interactions can have comparable frequencies. In online matching, we use the top 200 most recent behaviors. All graph-based models have the same online matching strategy. The batch sizes and the negative sample numbers of the intra-CL, inter-CL, and neighbor-similarity based losses are 4, 096 and 10. The temperature τ is set to be 1. For all models, we conduct a grid search to select parameters. Parameter analyses of CL loss weights are given in Sec. 5.7. All models share the same experimental settings and multi-domain behaviors for fair comparisons.

5.4 Evaluation of CDR in Matching (RQ1)

5.4.1 Evaluation Protocols. We evaluate on the few-shot and strict cold-start domains separately. All models select top N items from the overall corpora for each test instance. Following classical matching models [26, 35, 36], we utilize the top N hit rate (HIT@N) as our evaluation metric. To simulate the real-world matching systems, we concentrate on larger N as 50, 100, 200, and 500 (we retrieve top 500 items in online). We should double clarify that CCDR focuses on CDR in matching, which ***cares whether good items are retrieved***, not the specific ranks that should be measured by ranking. Hence, HIT@N is much more suitable for matching than ranking metrics such as AUC and NDCG. We also evaluate the diversity via a classical aggregate diversity metric named item coverage [8].

5.4.2 Experimental Results. From Table 2 we can observe that:

(1) CCDR achieves significant improvements over all baselines on all HIT@N in both two domains, with the significance level $p < 0.01$ (the deviations of CCDR are within ± 0.0004 in HIT@500). It indicates that CCDR can learn high-quality matching embeddings and well transfer useful knowledge to the target domain via CL. The improvements of CCDR mainly derive from three aspects: (a) The intra-CL enables more sufficient and balanced training via SSL with selected negative samples, which successfully alleviates the data sparsity and popularity bias issues. (b) The inter-CL builds interactions across different domains via three CL tasks, which multiplies the knowledge transfer via heterogeneous bridges. (c) The diversified preference network, CL losses, and multi-channel matching cooperate well with each other. The similarities used in online matching are directly optimized via losses in Eq. (9).

(2) CCDR has large improvement percentages on the challenging strict cold-start domain (55% improvement on HIT@500 over SSCDR+), where users have no behaviors on target items. It is natural since the combination of the diversified preference network and user/taxonomy/neighbor based inter-CL tasks can transfer more diversified preferences via more cross-domain bridges. Moreover, comparing with different CCDR models, we find that both intra-CL and inter-CL are effective, while inter-CL plays a more important role in CDR. We also find that CCDR has 4.2% and 6.0% improvements on the diversity metric *item coverage* [8] compared to the best-performing GraphDR+ in two domains. It indicates that CCDR has better performances on the diversity via CL tasks.

(3) Among baselines, we find that ICAN performs better in the few-shot domain, while SSCDR+ performs better in the strict cold-start domain. It is because that ICAN strongly relies on the feature

field interactions between behaviors in different domains, which are extremely sparse or even missing in the strict cold-start scenarios. In contrast, SSCDR+ benefits from cross-domain mapping. Moreover, classical matching methods such as GraphDR+ perform worse than CCDR. It implies that simply mixing behaviors in two domains may not get good performances, since the unbalanced domain data will confuse the user preference learning in the target domain.

(4) Comparing with different CDR methods, we observe that the CL-based methods are the most effective compared to knowledge distillation (e.g., cross-domain KD) and embedding mapping (e.g., EMCDR+ and SSCDR+). It is because that (a) contrastive learning can provide a sufficient and balanced training via SSL, and (b) CCDR conducts knowledge transfer via not only aligned users, but also taxonomies and neighbors. In this case, the popularity bias and data sparsity issues in the CDR part can be largely alleviated.

5.5 Online Evaluation (RQ2)

5.5.1 Evaluation Protocols. To verify the effectiveness of CCDR in real-world scenarios, we conduct an online A/B test on a well-known online recommendation system named WeChat Top Stories. Precisely, we deploy CCDR and several competitive baselines in the matching module of a relatively cold-start domain as in Sec. 4, with the ranking module unchanged. The online baseline is the GraphDR+ (target) model trained solely on the target domain. In online evaluation, we focus on three online metrics in the target domain: (1) CTR, (2) average user duration per capita, and (3) average share rate per capita. We conduct the A/B test for 8 days.

Model	CTR	Duration	Share rate
GraphDR+ (source+target)	+1.121%	+2.096%	+1.416%
Sub-graph CL	+2.182%	+2.213%	+2.378%
Cross-domain KD	+2.056%	+2.959%	+6.244%
CCDR (inter-CL)	+3.375%	+4.354%	+7.947%
CCDR (inter-CL+intra-CL)	**+14.368%**	**+6.623%**	**+10.401%**

Table 3: Online A/B tests on WeChat Top Stories.

5.5.2 Experimental Results. Table 3 shows the online improvement percentages of all models. We can find that:

(1) CCDR significantly outperforms all models in three metrics with the significance level $p < 0.01$. Note that all models are based on the same single-domain model in CCDR (i.e., GraphDR+). It reconfirms the effectiveness of the intra-domain and inter-domain contrastive learning. We jointly consider multiple behaviors such as click, share and like to build the diversified preference network, and use a neighbor-similarity based loss to learn user diverse preferences. Hence, CCDR has improvements on different metrics, which reflects user's real satisfaction more comprehensively.

(2) Comparing with the base model that only considers the target domain, we know that the source domain's information is essential. Looking into the differences among GraphDR+, Sub-graph CL (i.e., CCDR (intra-CL)), and CCDR (inter-CL), we can find that both intra-CL and inter-CL are effective in online scenarios. Moreover, CCDR models outperform GraphDR+ and Cross-domain KD, which also verifies the advantages of inter-CL over simple multi-domain mixing and cross-domain knowledge distillation in knowledge transfer.

5.6 Ablation Study (RQ3)

In this subsection, we further compare CCDR with its several ablation versions to show the effectiveness of different CL tasks. Table 4 displays the HIT@N results on both few-shot and strict cold-start domains. We find that: (1) Both intra-CL and inter-CL are essential in few-shot and cold-start domains. Inter-CL contributes the most to the CDR performances, since it is strongly related to the knowledge transfer task in CDR and fits well with the neighbor-similarity based loss of the single-domain model. (2) The intra-CL task also significantly improves the matching in CDR, while it just achieves slight improvements on the strict cold-start domain. The power of intra-CL will be multiplied when there are more user behaviors in the target domain. (3) From the second ablation group, we observe that all three inter-CL tasks can provide useful information for CDR. We observe that the user-based inter-CL functions well in the few-shot domain (since it has more user-related interactions), while taxonomy-based inter-CL achieves higher improvements in the cold-start domain. Note that CCDR does not use the user channel and adopt L_{inter_u} in the strict cold-start domain, since the cold-start user nodes are isolated in the target domain with no behaviors.

Model	few-shot		strict cold-start	
	HIT@50	HIT@500	HIT@50	HIT@500
CCDR	0.0225	0.1442	0.0161	0.0987
w/o intra-CL	0.0213	0.1327	0.0154	0.0931
w/o inter-CL	0.0174	0.1099	0.0118	0.0732
w/o L_{inter_u}	0.0211	0.1376	N/A	N/A
w/o L_{inter_t}	0.0219	0.1403	0.0149	0.0901
w/o L_{inter_n}	0.0221	0.1411	0.0157	0.0948

Table 4: Ablation tests on CDR-427M.

5.7 Model Analyses (RQ4)

We further conduct two model analyses on different weights of the intra-CL and inter-CL losses. Fig. 4 displays the HIT@500 results of different intra- and inter- CL weights (λ_3 and λ_4) on both few-shot and strict cold-start domains. We can find that: (1) as the loss weight increases, the HIT@500 results of both intra-CL and inter-CL losses first increase and then decrease. The best parameters are $\lambda_3 = 1.5, \lambda_4 = 0.6$. Note that the parameter analysis is carried out around the optimal parameter point. (2) The performance trends of two CL loss weights are relatively consistent on both few-shot and strict cold-start domains. Moreover, CCDR models with different CL loss weights still outperform all baselines, which verifies the robustness and usability of CCDR in real-world scenarios.

6 CONCLUSION AND FUTURE WORK

In this work, we propose a novel CCDR framework to deal with CDR in matching. We adopt the intra-CL to alleviate the data sparsity and popularity bias issues in matching, and design three inter-CL tasks to enable more diverse and effective knowledge transfer. CCDR achieves significant offline and online improvements on different scenarios, and is deployed on real-world systems. In the future, we

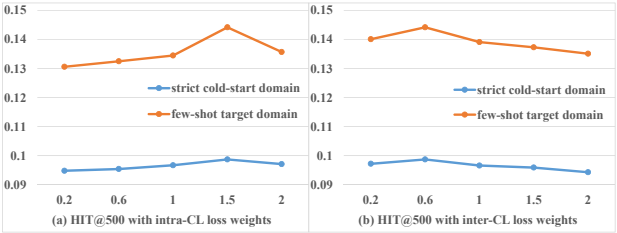

Figure 4: Results of different intra-/inter- CL loss weights.

will explore more sophisticated inter-CL tasks to further improve the effectiveness and diversity of knowledge transfer. We will also extend inter-CL to other CDR models and knowledge transfer tasks.

REFERENCES

[1] Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. Controllable multi-interest framework for recommendation. In *Proceedings of KDD*.

[2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *Proceedings of ICML*.

[3] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of RecSys*.

[4] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv:2006.07733* (2020).

[5] Xiaobo Hao, Yudan Liu, Ruobing Xie, Kaikai Ge, Linyao Tang, Xu Zhang, and Leyu Lin. 2021. Adversarial Feature Translation for Multi-domain Recommendation. In *Proceedings of KDD*.

[6] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of CVPR*.

[7] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of SIGIR*.

[8] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. 2004. Evaluating collaborative filtering recommender systems. *TOIS* (2004).

[9] Guangneng Hu, Yu Zhang, and Qiang Yang. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of CIKM*.

[10] Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based retrieval in facebook search. In *Proceedings of KDD*.

[11] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Trans. Big Data* (2019).

[12] SeongKu Kang, Junyoung Hwang, Dongha Lee, and Hwanjo Yu. 2019. Semi-supervised learning for cross-domain recommendation to cold-start users. In *Proceedings of CIKM*.

[13] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* (2009).

[14] Pan Li and Alexander Tuzhilin. 2020. DDTCDR: Deep dual transfer cross domain recommendation. In *Proceedings of WSDM*.

[15] Qi Liu, Ruobing Xie, Lei Chen, Shukai Liu, Ke Tu, Peng Cui, Bo Zhang, and Leyu Lin. 2020. Graph Neural Network for Tag Ranking in Tag-enhanced Video Recommendation. In *Proceedings of CIKM*.

[16] Fuyu Lv, Taiwei Jin, Changlong Yu, Fei Sun, Quan Lin, Keping Yang, and Wilfred Ng. 2019. SDM: Sequential deep matching model for online large-scale recommender system. In *Proceedings of CIKM*.

[17] Tong Man, Huawei Shen, Xiaolong Jin, and Xueqi Cheng. 2017. Cross-Domain Recommendation: An Embedding and Mapping Approach.. In *Proceedings of IJCAI*.

[18] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).

[19] Wentao Ouyang, Xiuwu Zhang, Lei Zhao, Jinmei Luo, Yu Zhang, Heng Zou, Zhaojie Liu, and Yanlong Du. 2020. MiNet: Mixed Interest Network for Cross-Domain Click-Through Rate Prediction. In *Proceedings of CIKM*.

[20] Arkadiusz Paterek. 2007. Improving regularized singular value decomposition for collaborative filtering. In *Proceedings of KDD cup and workshop*.

[21] Tieyun Qian, Yile Liang, Qing Li, and Hui Xiong. 2020. Attribute Graph Neural Networks for Strict Cold Start Recommendation. *TKDE* (2020).

[22] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. 2020. Gcc: Graph contrastive coding for graph neural network pre-training. In *Proceedings of KDD*.

[23] Steffen Rendle. 2010. Factorization machines. In *Proceedings of ICDM*.

[24] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. Fitnets: Hints for thin deep nets. In *Proceedings of ICLR*.

[25] Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. 2019. Autoint: Automatic feature interaction learning via self-attentive neural networks. In *Proceedings of CIKM*.

[26] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of CIKM*.

[27] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of ICLR*.

[28] Yinwei Wei, Xiang Wang, Qi Li, Liqiang Nie, Yan Li, Xuanping Li, and Tat-Seng Chua. 2021. Contrastive Learning for Cold-Start Recommendation. In *Proceedings of MM*.

[29] Ga Wu, Maksims Volkovs, Chee Loong Soon, Scott Sanner, and Himanshu Rai. 2019. Noise contrastive estimation for one-class collaborative filtering. In *Proceedings of SIGIR*.

[30] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised graph learning for recommendation. In *Proceedings of SIGIR*.

[31] Yiqing Wu, Ruobing Xie, Yongchun Zhu, Xiang Ao, Xin Chen, Xu Zhang, Fuzhen Zhuang, Leyu Lin, and Qing He. 2022. Multi-view Multi-behavior Contrastive Learning in Recommendation. In *Proceedings of DASFAA*.

[32] Yiqing Wu, Ruobing Xie, Yongchun Zhu, Fuzhen Zhuang, Xiang Ao, Xu Zhang, Leyu Lin, and Qing He. 2022. Selective Fairness in Recommendation via Prompts. In *Proceedings of SIGIR*.

[33] Yiqing Wu, Ruobing Xie, Yongchun Zhu, Fuzhen Zhuang, Xu Zhang, Leyu Lin, and Qing He. 2022. Personalized Prompts for Sequential Recommendation. *arXiv preprint arXiv:2205.09666* (2022).

[34] Chaojun Xiao, Ruobing Xie, Yuan Yao, Zhiyuan Liu, Maosong Sun, Xu Zhang, and Leyu Lin. 2021. UPRec: User-Aware Pre-training for Recommender Systems. *arXiv preprint arXiv:2102.10989* (2021).

[35] Ruobing Xie, Qi Liu, Shukai Liu, Ziwei Zhang, Peng Cui, Bo Zhang, and Leyu Lin. 2021. Improving Accuracy and Diversity in Matching of Recommendation with Diversified Preference Network. *IEEE Transactions on Big Data* (2021).

[36] Ruobing Xie, Zhijie Qiu, Jun Rao, Yi Liu, Bo Zhang, and Leyu Lin. 2020. Internal and Contextual Attention Network for Cold-start Multi-channel Matching in Recommendation. In *Proceedings of IJCAI*.

[37] Zhe Xie, Chengxuan Liu, Yichi Zhang, Hongtao Lu, Dong Wang, and Yue Ding. 2021. Adversarial and Contrastive Variational Autoencoder for Sequential Recommendation. In *Proceedings of WWW*.

[38] Jun Xu, Xiangnan He, and Hang Li. 2018. Deep learning for matching in search and recommendation. In *Proceedings of SIGIR*.

[39] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Proceedings of NeurIPS*.

[40] Junliang Yu, Hongzhi Yin, Jundong Li, Qinyong Wang, Nguyen Quoc Viet Hung, and Xiangliang Zhang. 2021. Self-Supervised Multi-Channel Hypergraph Convolutional Network for Social Recommendation. In *Proceedings of WWW*.

[41] Zheni Zeng, Chaojun Xiao, Yuan Yao, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2021. Knowledge transfer via pre-training for recommendation: A review and prospect. *Frontiers in big Data* (2021).

[42] Cheng Zhao, Chenliang Li, and Cong Fu. 2019. Cross-domain recommendation via preference propagation GraphNet. In *Proceedings of CIKM*.

[43] 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 SIGIR*.

[44] Chang Zhou, Jianxin Ma, Jianwei Zhang, Jingren Zhou, and Hongxia Yang. 2021. Contrastive Learning for Debiased Candidate Generation in Large-Scale Recommender Systems. In *Proceedings of KDD*.

[45] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *Proceedings of CIKM*.

[46] Han Zhu, Xiang Li, Pengye Zhang, Guozheng Li, Jie He, Han Li, and Kun Gai. 2018. Learning Tree-based Deep Model for Recommender Systems. In *Proceedings of KDD*.

[47] 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 SIGIR*.

[48] 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 WSDM*.

A APPENDIX

A.1 Online Deployment and Efficiency

A.1.1 Online Deployment. We have deployed CCDR on WeChat Top Stories (a well-known recommendation system widely used by millions of users) for more than six months. In practical systems, the quality, novelty, and diversity of item contents are the dominating factors that strongly impact user experiences. Hence, real-world recommendation systems often need to introduce new data sources as cold-start domains to add more contents and attract users. CCDR is deployed on the matching module to rapidly retrieve high-quality items. The relatively few-shot domain is viewed as the target domain, while other many-shot domains are combined as the source domain. We retrieve top 500 items for each target domain in matching. These retrieved items of CCDR are combined with candidates of other domains, and are then fed into the following ranking module for the final display. The computation cost is acceptable for the offline daily training in practical systems. CCDR is a straightforward, effective, and efficient implementation for CDR in matching, which can also be easily adopted in other new domains.

A.1.2 Online Efficiency. Real-world matching should deal with more than million-level candidates, where a linear prediction complexity w.r.t. the huge corpus size is unacceptable [46]. In CCDR, all embedding-based retrievals are conducted via the node embedding similarities learned in Eq. (9), which can be pre-calculated in offline as stated in Sec. 4. The top-100 nearest items for all heterogeneous nodes are also pre-indexed in offline for efficiency. Therefore, we do not need to calculate $sim(\cdot,\cdot)$ in online, and the online computation cost mainly locates in the indexing (finding top-100 nearest item embeddings via indexes) and sorting part (selecting the final top-500 items via s_i) of Eq. (10) to Eq. (12). We have an industry-level online indexer with cache technologies, which could largely accelerate the indexing part to nearly $O(1)$ and make sure it won't become a computation bottleneck. Since we have six channels and select top-100 nearest items for each heterogeneous object, the computation complexity of the indexing part is nearly $O(600n)$ (n is the length of the user historical behavior, and the maximum user behavior length n is 200). Similarly, the sorting part is $O(\log(600n))$. In conclusion, the overall online computation complexity is acceptable, and contrastive learning won't bring in additional online computation. As for the memory cost, CCDR only needs space to store the top-N nearest neighbor indexes for all nodes, which is acceptable for an industrial large-scale recommendation system.

A.2 Dataset

We build the CDR-427M dataset from a real-world multi-domain recommendation system in WeChat Top Stories. Specifically, we select two relatively new domains as the target domains, including a few-shot target domain where each user may have several behaviors, and a strict cold-start target domain where all users have no target-domain behaviors in the train set [21]. Users with the same account in different domains are aligned naturally. To simulate real-world scenarios and make full use of all information, other many-shot domains existing in our system (the main feed) are combined as the source domain. Hence, most users are overlapped in the global large source domain. Note that we evaluate on the strict cold-start setting to simulate the practical scenario when some new data sources are added to the existing recommendation system for the first time (which is not rare in the industry). The users are randomly selected and the dataset is split according to the chronological order. The item attributes used in CCDR help knowledge transfer, which can be obtained in most recommender systems. The inter-CL is universal. All data are preprocessed via data masking to protect the user's privacy.

We build the diversified preference networks for three domains separately following Sec. 3.2. The tags and categories are shared across different domains. Each item has one media (i.e., the item's producer). The words are extracted from the items' titles. These four types of item attributes widely exist and could be easily collected in most industrial recommendation systems. Besides user-item interactions, we further consider the session information to build item behavioral correlations. In total, there are 2.49, 0.5, 0.73 million nodes and 175, 29.8, 15.0 million edges in the diversified preference networks of the source domain, few-shot target domain, and strict cold-start domain, respectively. We bring in heterogeneous information to (a) better learn node representations in each domain besides sparse user-item interactions, and (b) make full use of all possible knowledge transfer paths between different domains. In fact, CCDR has been widely used in multiple online domains.

A.3 Motivations and Implementation Details

A.3.1 Motivations of Model Designs. In this work, we start from the challenging cold-start issue when adding new data sources or new recommendation scenarios, which often happens in practical recommendations. Cross-domain recommendation is a promising solution to this issue. However, it is trivial and time-consuming to design effective customized domain-specific models for different target domains. Hence, our proposed CDR model should not only be **effective** and **efficient**, but also be **universal** and **easy to deploy**. We should double clarify that the *online performance* is the central objective. Therefore, rather than fancy complicated methods, we prefer **simple** and effective models according to our industry taste.

Under these principles, we first bring in contrastive learning to improve representation learning and knowledge transfer in CDR with different intra-CL and inter-CL tasks, which are two of our main contributions. We bring CL into CDR of matching, since CL-based self-supervised training enables more sufficient training especially for those long-tail and unaligned items in the target domain. We do not propose new CL-based architectures or design fancy complicated GNN aggregators, since we find that the classical InfoNCE and GAT-based aggregation have already achieved significant offline and online improvements in different domains. Instead, we propose new types of augmentations for CDR, which are also an essential component in CL. For GNN aggregator, we directly use GAT, since we suggest that all heterogeneous embeddings are in the same semantic space, as we should retrieve diversified items via similarities between different heterogeneous objects in online multi-channel matching. The single-domain GNN modeling is not the main focus of this work. Nevertheless, it is still convenient to conduct other heterogeneous GNN models in node aggregation. Although more sophisticated CL tasks and GNN models may further improve the performance, the universality, offline/online efficiency,

and deployment costs will be worse. How to balance the four real-world principles with enhanced models could be our future work. For the construction of the diversified preference network, we follow the existing model's setting [35] and modify it according to our system's characteristics. It is also fine to use other informative interactions for more sufficient training.

As for the specific advantages and motivations of different components, the advantages of CCDR are introduced in Sec. 1 and Sec. 5.4 with evaluation results. The advantages of the diversified preference network and the information brought by different edges are given in Sec. 3.2. The motivations of both intra-CL and inter-CL are in Sec. 3.4, Sec. 3.5, and Sec. 5.6. The feasibility and advantages of the cross-domain multi-channel matching are also concluded in Sec. 4. Different components and model designs cooperate for the same goals of improving representation learning and cross-domain knowledge transfer, aiming to alleviate the sparsity issues in CDR.

A.3.2 Implementation Details. Sec. 5.3 has shown most essential hyper-parameters of CCDR. For the basic parameters such as the input and output dimensions of heterogeneous node embeddings and the numbers of weighted neighbor sampling, we empirically set them according to our existing online GNN-based matching models to balance effectiveness and efficiency under the practical computation limitation. We conduct a grid search to select sensitive parameters. The most essential parameters are the weights of different losses. For the CL loss weights, we have chosen different intra-CL and inter-CL loss weights among $\{0.2, 0.6, 1.0, 1.5, 2.0\}$, and find the best parameters as $\lambda_1, \lambda_2, \lambda_3, \lambda_4 = 1.0, 1.0, 1.5, 0.6$. Fig. 4 shows the results of different loss weights. For the possible channel weights in online cross-domain multi-channel matching, we empirically use the online default setting (i.e., equal channel weights). It is also easy to adjust the heterogeneous channels' weights for practical demands (e.g., if we want to highlight user's media-level preferences, we can choose bigger media channel weights), which is not the main concern of this work. For the batch size, we have tested among $\{128, 512, 1024, 4096\}$, and select the best $4,096$ with larger batch size [2].

A.4 Evaluation Details

A.4.1 Baselines. We focus on cross-domain recommendation in matching. Therefore, we conduct three groups of competitive baselines for comprehensive comparisons. The first group is the classical matching methods, including three classical industrial matching models (i.e., FM, AutoInt, and GraphDR+). Note that both the source and target domains' behaviors are considered in these baselines for fair comparisons. GraphDR [35] is the SOTA graph-based matching model. We enhance the original GraphDR model for the CDR task as GraphDR+, which also acts as the single-domain model of our CCDR for fair comparisons. We do not compare with other matching models like ComiRec [1] and TDM [46], for they are designed for different challenges and are nontrivial to be modified for CDR. The second group is cross-domain/multi-domain methods, including EMCDR, SSCDR, and ICAN. For EMCDR and SSCDR, we also adopt GraphDR+ as the single-domain model. ICAN is the SOTA multi-domain matching model. We also do not compare with CoNet [9] and CATN [43], for they are designed for ranking and will be extremely time-consuming in matching. The third group contains another knowledge transfer method of knowledge distillation, which is widely used in practice. Moreover, existing graph-based intra-CL techniques [22, 30] are also implemented as baselines and ablation versions.

A.4.2 Offline Evaluation Metrics. We concentrate on the CDR in matching. In practical large-scale recommendation systems, there are usually million-level item candidates, where complicated user-item feature interaction modeling is hard to be calculated on the overall corpus. Hence, the two-step recommendation system is proposed [3], where the matching module focuses more on efficiently retrieving a small subset of item candidates from the million-level large corpora, while the ranking module aims to give the specific top-k item ranks accurately. In our online system, we select top-500 items in the matching module for the following ranking. The specific item ranks of these top-500 items are not important, since they will be re-ranked in the following ranking and mixing modules after matching. Hence, following classical matching models [26, 35, 36], we use the top N hit rate (HIT@N) as our evaluation metric instead of ranking metrics such as NDCG and AUC. To simulate the real-world matching systems, we report the HIT@N with larger N as 50, 100, 200, and 500. Moreover, since matching often struggles with the popularity issue, we also evaluate the recommendation diversity via a classical aggregate diversity metric named item coverage [8]. The item coverage measures the percentage of a dataset that the recommendation model is able to provide predictions for in the test set. In conclusion, we find significant improvements in both recommendation accuracy and diversity. Note that although we do not explicitly optimize diversity, CCDR still has 4.2% and 6.0% improvements on item coverage over GraphDR+.

A.4.3 Online Evaluation Metrics. In online evaluation, we focus on three representative online metrics in the target domain: (1) Click-through-rate (CTR), (2) average user duration per capita (Duration), and (3) average share rate per capita (Share rate). We have:

$$\text{CTR} = \frac{\text{\# of click counts}}{\text{\# of impression counts}}, \qquad \text{Duration} = \frac{\text{all duration}}{\text{\# of users}},$$
$$\text{Share rate} = \frac{\text{\# of share counts}}{\text{\# of users}}. \tag{13}$$

CTR measures the click-related recommendation accuracy. Duration focuses on the dwell time, which can better reflect user's real satisfaction. The share rate can reflect user's satisfaction via a more high-cost user social behavior. We conduct the A/B test for 8 days.