

Pre-trained Recommender Systems: A Causal Debiasing Perspective

Ziqian Lin*[†]
zlin284@wisc.edu
University of Wisconsin-Madison
USA

Hao Ding*
haodin@amazon.com
AWS AI Labs
USA

Nghia Trong Hoang[‡]
trongnghia.hoang@wsu.edu
Washington State University
USA

Branislav Kveton
bkveton@amazon.com
AWS AI Labs
USA

Anoop Deoras
adeoras@amazon.com
AWS AI Labs
USA

Hao Wang
howngz@amazon.com
AWS AI Labs
USA

ABSTRACT

Recent studies on pre-trained vision/language models have demonstrated the practical benefit of a new, promising solution-building paradigm in AI where models can be pre-trained on broad data describing a generic task space and then adapted successfully to solve a wide range of downstream tasks, even when training data is severely limited (e.g., in zero- or few-shot learning scenarios). Inspired by such progress, we investigate in this paper the possibilities and challenges of adapting such a paradigm to recommender systems. In particular, we develop a generic recommender that captures universal interaction patterns by training on generic user-item interaction data extracted from different domains, which can then be fast adapted to improve few-shot learning performance in unseen new domains (with limited data).

However, unlike vision/language data which share strong conformity in the semantic space, universal patterns underlying recommendation data collected across different domains (e.g., different countries) are often occluded by both in-domain and cross-domain biases implicitly imposed by the cultural differences in their user and item bases. Such heterogeneous biases tend to hinder the effectiveness of pre-trained models. To address this challenge, we further introduce and formalize a causal debiasing perspective, which is substantiated via a hierarchical Bayesian deep learning model, named **PREREC**. Our empirical studies on real-world data show that the proposed model could significantly improve the recommendation performance in zero- and few-shot learning settings under both cross-market and cross-platform scenarios. Our code is publicly available at GitHub for reproducibility¹.

*Both authors contributed equally to this research.

[†]Work done during the author's internship at AWS AI Labs.

[‡]Work done at AWS AI Labs.

¹<https://github.com/myhakureimu/PreRec>

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.

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

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0371-3/24/03...\$15.00

<https://doi.org/10.1145/3616855.3635779>

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender Systems, Pre-Trained Models, Causality, Probabilistic Methods, Bayesian Inference

ACM Reference Format:

Ziqian Lin, Hao Ding, Nghia Trong Hoang, Branislav Kveton, Anoop Deoras, and Hao Wang. 2024. Pre-trained Recommender Systems: A Causal Debiasing Perspective. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, Mexico, USA, 13 pages. <https://doi.org/10.1145/3616855.3635779>

1 INTRODUCTION

Recommender systems (RecSys) [7, 10, 13, 18, 52] recommend relevant items from a large inventory based on the personal interests of the user, which mitigates the information overload issue on the Internet. They have been proven to be effective on a wide range of application scenarios including e-commerce and online media. An essential problem for RecSys is data scarcity, which restrains the recommendation performance. Recent years have witnessed the great success of the pre-trained language model (PLM) such as BERT [6] and GPT [3] in natural language processing (NLP), which improves the performance on the target domain by leveraging massive training data from other domains. A natural question that arises is whether it is feasible to build a pre-trained RecSys that is generally applicable across different domains.

Traditional approaches, including collaborative filtering-based methods [31, 40] and session-based methods [13, 18, 29], typically learn item embeddings indexed by domain-specific item IDs (also known as *item ID embeddings*) through interaction data. Such item ID embeddings are transductive and are not generalizable to unseen items, which causes the in-domain item cold-start issue and also becomes a blocker for transferring knowledge from one domain to another. Of course we can infer the embedding of the unseen item based on its metadata, however different domains are likely to possess different sets of item metadata. For example, the item metadata in domain m is $\{director, actor\}$ while in domain n is $\{genre, rating\}$, the knowledge that the model learned based on item metadata in domain m can never be transferred to domain n and vice versa.

A new line of deep learning recommenders, pioneered by collaborative deep learning (CDL) [40] and its variants [23, 39], seamlessly incorporate content information into deep recommenders, thereby opening up the possibility of pretraining-based recommenders and significantly alleviating data scarcity (and cold-start) problems in recommender systems. Inspired by CDL, we argue that the key to learning a generalizable model lies in capturing knowledge grounded on a universal feature space. In order to bridge the gap between different domains, ZESRec [8] first proposed to use generic item textual description (e.g., product description, movie synopsis, and news content) to produce the item universal embedding, while the user universal embedding is computed via a sequential model (Gated Recurrent Units or GRU, Transformer, etc.) aggregating item universal embeddings for items in the user history. Despite the promising results, ZESRec limits itself to conducting pre-training on a single source domain and inference on a single target domain. As follow-up work, UniSRec [15] further extends support for multi-domain pre-training by designing cross-domain contrastive learning tasks, and evaluate the pre-trained model on multiple target domains. However, UniSRec fails to consider the bias either within each domain or across domains during pre-training, while both types of bias may lead to drift in user interests, item properties, and user behavioral patterns. Training the model on signals mixed with bias and general knowledge without considerations of debiasing will lead to overfitting and impairs the generalizability of the model, as shown in Sec. 4.

In this work, we aim to design a universally generalizable recommender that can be pre-trained on multiple source domains and fine-tuned on different target domains. We start by identifying two types of bias:

- (1) **In-domain bias** considers the biases that take effect within each domain, and one example is the popularity bias that affects not only the item exposure rate, but also the user behavioral pattern due to users tend to follow the majority and interact with trending items (user conformity effect [28]).
- (2) **Cross-domain bias** considers the item bias introduced by the unique domain properties, for instance, a domain launches a promotion campaign for a set of items which influences both item price and user behavioral patterns. Besides, each domain has its characters, which will form its distinctive user community; this causes the shift in user interests across different domains.

To address the aforementioned challenges, we propose a novel Bayesian deep learning framework, dubbed Pre-trained Recommender Systems (PREREC), which is equipped with a causal debiasing mechanism and distills general knowledge from massive multi-domain data. Our model isolates the effects of both in-domain bias and cross-domain bias by explicitly incorporating corresponding bias terms into the modeling during pre-training, and performs debiasing by (1) enforcing causal intervention and neutralizing the cross-domain bias and (2) inferring and removing the in-domain bias during inference in the target domain. In this way, our model utilizes the bias terms to explain away the variability originating from in-domain and cross-domain dynamics, and is capable of capturing universal item properties and then inferring universal user

interests accordingly. In order to adapt to the new domain, PRE-REC is enabled to capture the bias in the target domain during fine-tuning and integrate such bias during inference to improve recommendation performance. To summarize our contributions:

- We first identify the in-domain bias and the cross-domain bias existing in the multi-domain data which may potentially undermine the generalizability of the pre-trained RecSys.
- We design a novel Bayesian deep learning model, named PRE-REC, which distills general knowledge from multi-domain data. It captures the in-domain bias and the cross-domain bias, and adds casual intervention for debiasing.
- We extensively conduct experiments on datasets collected from different domains and evaluate PRE-REC on multiple target domains in three settings including zero-shot, incremental training, and fine-tuning, under both cross-market and cross-platform scenarios. The empirical results demonstrate the effectiveness of our model.

2 PRELIMINARY

In this paper, we focus on training a generalizable recommender system on multiple source domains; to achieve this, the model needs item universal representations and user universal representations. As preliminaries, here we introduce the item universal embedding network (item UEN) and the user universal embedding network (user UEN).

Item Universal Embedding Network. An item universal embedding network, denoted as $f_e(\cdot)$, generates item j 's embedding based on its textual description \mathbf{X}_j . The network consists of a pre-trained language model (PLM) to extract generic semantic features from \mathbf{X}_j ; in our work we adopt a multi-lingual version of Sentence-BERT² [33], denoted as $f_{BERT}(\cdot)$, followed by a single-layer neural network $f_{NN}(\cdot)$. The tokenizer of Sentence-BERT ingests \mathbf{X}_j to yield a set of tokens $\{t_1, t_2, \dots, t_T\}$ before sending the input to the model, formally:

$$\mathbf{m}_j = f_e(\mathbf{X}_j) = f_{NN}(f_{BERT}(\{t_1, t_2, \dots, t_T\})), \quad (1)$$

where \mathbf{m}_j represents item universal embedding of item j .

User Universal Embedding Network. We denote the item universal embeddings (in chronological order) of historical items for user i as an embedding matrix $\mathbf{H}_i \in \mathbb{R}^{N_u \times B}$, where N_u and B represent the number of items in the history of user i and the hidden dimension of item universal embedding. Note that \mathbf{H}_i is time-sensitive, and we ignore the time indexing for simplicity. A user universal embedding network is essentially a sequential model serving as an aggregation function over \mathbf{H}_i :

$$\mathbf{u}_i = f_{seq}(\mathbf{H}_i), \quad (2)$$

where \mathbf{u}_i stands for user universal embedding of user i , $f_{seq}(\cdot)$ denotes any type of sequential model, and in this work, we consider a transformer-based recommender [18] due to its superior performance.

An existing issue for user UEN is that it cannot generate user embedding for users without interactions. To address this issue, one can pad a learnable dummy item at the front of each user sequence,

²<https://huggingface.co/sentence-transformers/stsb-xlm-r-multilingual>

which can be seen as a prior of user behavioral patterns. In our pre-training stage, we consider user communities in multiple domains, where each domain needs its own prior. We therefore adjust the user UEN to accommodate domain prior:

$$\mathbf{n}_i = f_{seq}(\mathbf{D}_k, \mathbf{H}_i), \quad (3)$$

where $\mathbf{D}_k \in \mathbb{R}^B$ denote the prior of domain k , \mathbf{n}_i denotes and the user embedding of user i . We assume \mathbf{D}_k is drawn from a zero-mean isotropic multivariate Gaussian distribution.

3 PRE-TRAINED RECOMMENDER SYSTEMS

In this section, we introduce PREREC, a flexible hierarchical Bayesian deep learning framework [39, 41, 42] and can be easily extended to any sequential model (in this work we adopt Transformer). It can be pre-trained on multiple source domains to distill knowledge grounded on a universal feature space; in this work we focus on the textual feature space as a possible instantiation of the universal feature space, but note that our model is generally applicable to any other modalities such as images. In general, PREREC works in three stages:

- (1) **Multi-domain Pre-training:** PREREC is pre-trained on the data collected from multiple source domains, taking into account both cross-domain and in-domain bias.
- (2) **Zero-shot Recommendation:** PREREC recommends in the target domain without fine-tuning on any target domain interactions.
- (3) **Fine-tuning:** PREREC is further fine-tuned on the retrieved target domain data, capturing domain bias and popularity bias in the target domain, and adjusting target domain inference accordingly.

3.1 Model Overview

Fig. 1 shows the graphical model for PREREC. (See the supplement for the model architecture from a neural network perspective.) Below we explain its rationale in detail:

- Variable $\mathbf{D}_k \in \mathbb{R}^B$ represents the (latent) properties of domain k such as the user community, promotion campaigns, as well as website design.
- Variables $\mathbf{U}_i \in \mathbb{R}^B$ and $\mathbf{H}_i \in \mathbb{R}^{N_u \times B}$ represent the genuine interests of user i and user i 's interaction history, respectively. Note that we ignore the time indexing of \mathbf{H}_i for simplicity.
- Variable $\mathbf{F}_j \in \mathbb{R}^C$ represents the popularity factors of item j , including three prominent factors: (1) number of interactions of item j , (2) number of interactions of all items (traffic volume), and (3) time. We ignore the time indexing of \mathbf{F}_j for simplicity.
- Variables $\mathbf{X}_j, \mathbf{Z}_j \in \mathbb{R}^B$, and $\mathbf{V}_j \in \mathbb{R}^B$ represent item j 's textual description, item j 's popularity properties derived from \mathbf{F}_j , and item j 's overall properties, respectively. Since a user only interacts with displayed items, we assume item j is exposed.
- Variable $\mathbf{R}_{ijk} \in \{0, 1\}$ is the interaction label denoting whether user i interacted with item j in domain k .
- Edges $\mathbf{D}_k \rightarrow \{\mathbf{U}_i, \mathbf{V}_j, \mathbf{R}_{ijk}\}$: Domain properties may influence user interests, item properties, and user behavioral patterns.
- Edge $\mathbf{F}_j \rightarrow \mathbf{Z}_j$: Item popularity factors decide the popularity properties.




Figure 1: The probabilistic graphical model (PGM) for our PREREC. \mathbf{U}_i and \mathbf{H}_i represent user i and corresponding user history; $\mathbf{V}_j, \mathbf{X}_j$, and \mathbf{Z}_j represent item j , its textual description (e.g., movie synopsis), and its popularity effect; \mathbf{F}_j represents all the prominent factors impact item popularity; \mathbf{D}_k represents domain k . K_s represents all the source domains, while J_k, I_k represent all the items and all the users in domain k , respectively. $\lambda_u, \lambda_v, \lambda_d, \lambda_z$ are hyperparameters related to distribution variance.

- Edge $\mathbf{H}_i \rightarrow \mathbf{U}_i$: What the user has interacted in the past may affect user's next move (e.g., a user who purchased a cell phone may want to purchase its accessories next).
- Edges $\{\mathbf{X}_j, \mathbf{Z}_j\} \rightarrow \mathbf{V}_j$: Item textual descriptions affect item properties and the popularity factors affect item exposure rate.
- Edge $\{\mathbf{U}_i, \mathbf{V}_j, \mathbf{D}_k, \mathbf{Z}_j\} \rightarrow \mathbf{R}_{ijk}$: Interaction depends on user interests \mathbf{U}_i , item properties \mathbf{V}_j , domain properties \mathbf{D}_k , and the user conformity effect caused by \mathbf{Z}_j .

Here $\mathbf{D}_k, \mathbf{U}_i, \mathbf{V}_j$, and \mathbf{Z}_j have the same hidden dimension, and we name them *latent domain embedding*, *latent user embedding*, *latent item embedding*, and *latent popularity embedding*, respectively. We also name \mathbf{X}_j as *item textual description*, and \mathbf{Z}_j as *popularity factors*. The corresponding conditional probabilities in the PGM are listed in Eqn. 6. In general, this is a hierarchical Bayesian deep learning (BDL) model [41, 42] with the Transformer as the deep component and the probabilistic graphical model in Fig. 1 as the graphical component to capture the causal relation.

Generative Process. Inspired by CDL [40] and ZESRec [8], the generative process in Fig. 1 for each domain k in source domains \mathcal{D} is defined as follows:

- (1) Draw a latent domain embedding $\mathbf{D}_k \sim \mathcal{N}(\mathbf{0}, \lambda_d^{-1} \mathbf{I}_B)$.
- (2) For each item j :
 - Draw a latent item offset vector $\epsilon_j^v \sim \mathcal{N}(\mathbf{0}, \lambda_v^{-1} \mathbf{I}_B)$.
 - Obtain the item universal embedding: $\mathbf{m}_j = f_e(\mathbf{X}_j)$.
 - Draw a latent popularity offset $\epsilon_j^z \sim \mathcal{N}(\mathbf{0}, \lambda_z^{-1} \mathbf{I}_B)$.
 - Compute the latent popularity embedding based on popularity factors \mathbf{F}_j directly derived from the data and the latent popularity offset vector: $\mathbf{Z}_j = f_{pop}(\mathbf{F}_j) + \epsilon_j^z$.

- Compute the latent item embedding: $\mathbf{V}_j = \epsilon_j^v + f_{item}(\mathbf{D}_k, \mathbf{Z}_j, \mathbf{m}_j)$, where $f_{item}(\cdot)$ represents the item encoder: $f_{item}(\mathbf{D}_k, \mathbf{Z}_j, \mathbf{m}_j) = \mathbf{D}_k + \mathbf{Z}_j + \mathbf{m}_j$.
- (3) For each user i :
- Draw a latent user offset vector $\epsilon_i^u \sim \mathcal{N}(\mathbf{0}, \lambda_u^{-1} \mathbf{I}_B)$.
 - Obtain the user embedding: $\mathbf{n}_i = f_{seq}(\mathbf{D}_k, \mathbf{H}_i)$.
 - Compute the latent user embedding: $\mathbf{U}_i = \epsilon_i^u + \mathbf{n}_i$.
 - Compute the recommendation score S_{ijk} for each tuple (i, j, k) , $S_{ijk} = f_{softmax}(\mathbf{U}_i^\top \mathbf{V}_j + \mathbf{D}_k \mathbf{W}_d + \mathbf{Z}_j \mathbf{W}_z)$, where \mathbf{W}_d and \mathbf{W}_z are trainable parameters, and for user i : $\mathbf{R}_{i*k} \sim \text{Cat}([\mathbf{S}_{ijk}]_{j=1}^{J_k})$, where $\mathbf{R}_{i*k} \in \{0, 1\}^{J_k}$ is a one-hot vector denoting an item ID ($\mathbf{R}_{ijk} = 1$ if j is the sampled item index from $\text{Cat}([\mathbf{S}_{ijk}]_{j=1}^{J_k})$, and $\mathbf{R}_{ijk} = 0$ otherwise); J_k denotes number of items in domain k , $*$ represents the collection of all elements in a specific dimension, $f_{softmax}(\cdot)$ represents the softmax function, and $\text{Cat}(\cdot)$ denotes a categorical distribution.

Below we make a few remarks on this generative process.

- \mathbf{D}_k , \mathbf{Z}_j , and \mathbf{m}_j in $f_{item}(\mathbf{D}_k, \mathbf{Z}_j, \mathbf{m}_j)$ capture cross-domain bias, in-domain (popularity) bias, and item content (e.g., text description), respectively.
- The latent offset vectors ϵ_j^v , ϵ_j^z , ϵ_i^u accommodate variances in item properties, popularity properties, and user interests. For example, ϵ_j^v provides flexibility that the final item vector \mathbf{v}_j can deviate from $f_{item}(\mathbf{D}_k, \mathbf{Z}_j, \mathbf{m}_j)$ and that ϵ_j^v can be different for different item j .
- For zero-shot recommendation (see Sec. 3.3), we remove all the latent offset vectors and set $\mathbf{D}_k = \mathbf{0}$ to eliminate noises.

Computing Popularity Factors \mathbf{F}_j . We now provide the details of computing the popularity factors \mathbf{F}_j and the popularity properties \mathbf{Z}_j for item j . Note that \mathbf{F}_j is time-sensitive, and we ignore the time indexing for simplicity. We divide the interaction data \mathbf{R} into a series of equal size time intervals $\mathcal{T} = \{T_1, T_2, \dots, T_{\mathcal{T}}\}$. Assuming time t , item j , and domain k are given, where t falls into the time interval T_{l+1} , we first obtain the number of interactions for item j in the former time interval T_l (prevent temporal leakage) and denote it as $c_j^{T_l}$, and use J_k to represent items in domain k . The popularity factors \mathbf{F}_j is computed as:

$$\mathbf{F}_j = [c_j^{T_1}/s_w^{T_1}, c_j^{T_2}/s_w^{T_2}, \dots, c_j^{T_{\mathcal{T}}}/s_w^{T_{\mathcal{T}}}],$$

where $s_w^{T_l}$ is a normalization term calculated as:

$$s_w^{T_l} = \left(\sum_{j \in J_k} (c_j^{T_l})^w / |J_k| \right)^{\frac{1}{w}}$$

We then use the trainable single layer neural network $f_{pop}(\cdot)$ to obtain $\mathbf{Z}_j = f_{pop}(\mathbf{F}_j)$. Note that the $f_{pop}(\cdot)$ is *generally applicable across different domains*. See the supplement for the design philosophy and visualization of computing popularity properties.

3.2 Multi-domain Pre-training

For multi-domain pre-training, our unique challenge is to leverage training data from different source domains *without introducing bias*. The sources of bias can be broadly summarized into two categories: (1) in-domain bias caused by variances within each domain such as popularity bias [29, 53, 57], and (2) cross-domain bias originated

from dynamics across different domains, including changes in user community, item catalogue, promotion campaigns, etc.

In-domain Bias and Explicit Confounders. For in-domain bias, we consider popularity bias \mathbf{Z}_j . We postulate that the popularity bias of item j is correlated with its number of interactions, traffic volume, and time; we use \mathbf{F}_j to represent all aforementioned popularity factors and hope to learn the mapping function from \mathbf{F}_j to \mathbf{Z}_j . We identify \mathbf{Z}_j as a *confounder* [32] which affects both exposed item \mathbf{V}_j and observed interactions \mathbf{R}_{ijk} .

Note that our PREREC can easily incorporate other types of in-domain bias which influence both exposed variables (users or items) and observed interactions, e.g. position bias [17], by adding corresponding confounders in the graphical model similar to \mathbf{Z}_j . Here we explicitly know that \mathbf{Z}_j models the popularity effect (since it is computed from the popularity factors); we therefore also call \mathbf{Z}_j the *explicit confounder*. This is in contrast with the cross-domain bias to be introduced next.

Cross-domain Bias and Implicit Confounders. Cross-domain bias such as changes in user community (US users' preferences may be very different from UK users') tends to be implicit or latent since relevant metadata is usually unavailable and therefore needs to be inferred from data.

To address this challenge, PREREC defines the domain property \mathbf{D}_k as an independent latent variable. \mathbf{D}_k affects user behavioral pattern, item properties, and user-item interaction; we therefore posit treat it as a causal *confounder* between user \mathbf{U}_i and interaction \mathbf{R}_{ijk} as well as between item \mathbf{V}_j and interaction \mathbf{R}_{ijk} . Here we implicitly model the domain property \mathbf{D}_k as a learnable latent variable which is not conditioned on any factors, we therefore call \mathbf{D}_k the *implicit confounder*.

Distinguishing between In- and Cross-domain Biases. Note that all users and items within each domain k share the same cross-domain bias induced by \mathbf{D}_k . In contrast, each item j in a domain has *different individual* in-domain bias induced by \mathbf{Z}_j . Therefore, during learning, PREREC will automatically extract the cross-domain bias shared across users and items into \mathbf{D}_k , while extracting the individual in-domain bias for each item j into \mathbf{Z}_j .

Pre-training. For the pre-training stage of PREREC, we assume the user-item interactions \mathbf{R}_{ijk} and item textual description \mathbf{X}_j are observed for each domain, while the popularity factors \mathbf{F}_j and user history \mathbf{H}_i can be directly derived from the interaction \mathbf{R}_{ijk} . With hyperparameters λ_u , λ_v , λ_d , and λ_z and given the graphical model Fig. 1, the maximum a posteriori (MAP) estimation on latent variables $(\mathbf{U}_i, \mathbf{V}_j, \mathbf{D}_k, \mathbf{Z}_j)$ can be decomposed as following:

$$\begin{aligned} & P(\mathbf{U}_i, \mathbf{V}_j, \mathbf{D}_k, \mathbf{Z}_j | \mathbf{R}_{ijk}, \mathbf{H}_i, \mathbf{X}_j, \mathbf{F}_j, \lambda_u, \lambda_v, \lambda_d, \lambda_z) \propto \\ & P(\mathbf{R}_{ijk} | \mathbf{U}_i, \mathbf{V}_j, \mathbf{D}_k, \mathbf{Z}_j) \cdot P(\mathbf{U}_i | \mathbf{H}_i, \mathbf{D}_k, \lambda_u) \cdot P(\mathbf{V}_j | \mathbf{X}_j, \mathbf{D}_k, \mathbf{Z}_j, \lambda_v) \\ & \quad \cdot P(\mathbf{Z}_j | \mathbf{F}_j, \lambda_z) \cdot P(\mathbf{D}_k | \lambda_d). \end{aligned} \quad (4)$$

We define the conditional probability over the observed interactions as:

$$P(\mathbf{R}_{ijk} | \mathbf{U}_i, \mathbf{V}_j, \mathbf{D}_k, \mathbf{Z}_j) = f_{softmax}(\mathbf{U}_i^\top \mathbf{V}_j + \mathbf{D}_k \mathbf{W}_d + \mathbf{Z}_j \mathbf{W}_z). \quad (5)$$




Figure 2: For zero-shot recommendation in the target domain, we perform causal intervention on U_i , V_j , and Z_j , i.e., $do(U_i = f_{seq}(D_k = 0, H_i))$, $do(V_j = f_e(D_k = 0, Z_j) = f_{pop}(F_j), f_{BERT}(X_j))$, and $do(Z_j = f_{pop}(F_j))$, to remove the cross-domain bias while injecting the in-domain bias in the target domain.

We assume Gaussian distributions on all latent variables in Fig. 1, i.e., U_i , V_j , D_k , and Z_j , as follows:

$$\begin{aligned} P(U_i | H_i, D_k, \lambda_u) &= \mathcal{N}(U_i; f_{seq}(D_k, H_i), \lambda_u^{-1} \mathbf{I}_B), \\ P(V_j | X_j, D_k, Z_j, \lambda_v) &= \mathcal{N}(V_j; f_{item}(D_k, Z_j, f_e(X_j)), \lambda_v^{-1} \mathbf{I}_B), \\ P(Z_j | F_j, \lambda_z) &= \mathcal{N}(Z_j; f_{pop}(F_j), \lambda_z^{-1} \mathbf{I}_B), \\ P(D_k | \lambda_d) &= \mathcal{N}(D_k; \mathbf{0}, \lambda_d^{-1} \mathbf{I}_B), \end{aligned} \quad (6)$$

where $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \lambda^{-1} \mathbf{I}_B)$ denotes the probability density function (PDF) of a Gaussian distribution with mean $\boldsymbol{\mu}$ and diagonal covariance $\lambda^{-1} \mathbf{I}_B$ for the variable \mathbf{x} ; $f_e(\cdot)$ and $f_{pop}(\cdot)$ are the learnable encoding functions for item embedding V_j and popularity embedding Z_j , respectively. In our work, we choose to adopt multi-layer perceptron (MLP) for both $f_e(\cdot)$ and $f_{pop}(\cdot)$.

As in CDL [40] and ZESRec [8], maximizing the posterior probability is equivalent to minimizing the negative log likelihood (NLL) of U_i , V_j , D_k and Z_j given R_{ijk} , X_j , H_i , F_j , λ_u , λ_v , λ_d and λ_z :

$$\begin{aligned} \mathcal{L} &= \sum_{k=1}^{K_s} \sum_{i=1}^{I_k} \sum_{j=1}^{J_k} -\log(f_{softmax}(U_i^T V_j + D_k W_d + Z_j W_z)) \\ &+ \frac{\lambda_z}{2} \sum_{j=1}^{J_k} \|Z_j - f_{pop}(F_j)\|^2 + \frac{\lambda_v}{2} \sum_{k=1}^{K_s} \sum_{j=1}^{J_k} \|V_j - f_e(D_k, Z_j)\|^2 \\ &+ \frac{\lambda_u}{2} \sum_{k=1}^{K_s} \sum_{i=1}^{I_k} \|U_i - f_{seq}(D_k, H_i)\|^2 + \frac{\lambda_d}{2} \sum_{k=1}^{K_s} \|D_k\|^2, \end{aligned} \quad (7)$$

where $W_d \in \mathbb{R}^B$, $W_z \in \mathbb{R}^B$ are learnable vectors. K_s denotes a set of source domains; for domain k , we have I_k users and J_k items.

3.3 Zero-shot Recommendation

Next we discuss how to use the pre-trained model under the zero-shot setting as defined in [8], which simulates an online environment where interaction data in the new domain is available only

during inference, and all new users and items are unseen in source domains.

Causal Zero-shot Recommendation. Given the pre-trained model, we enforce intervention on U_i , V_j and Z_j by performing the *do*-calculus [32] to eliminate the cross-domain bias while incorporate the in-domain bias in the target domain, as shown in Fig. 2. Specifically, we do not assume any domain properties and thus the posterior of D_k collapses to the prior $\mathbf{0}$; the popularity factors F_j are based on data statistics which can be derived on the fly in an online environment. Therefore, we compute $U_i = f_{seq}(D_k = \mathbf{0}, H_i)$ and $V_j = f_e(D_k = \mathbf{0}, Z_j = f_{pop}(F_j), f_{BERT}(X_j))$. Note that here we set $D_k = \mathbf{0}$ for U_i and V_j . This serves as an approximation to the output expectation over the distribution of D_k as input, and we found this approach achieved similar performance in practice. Following the *back-door formula* [32] we have:

$$\begin{aligned} &P(R_{ijk} | do(U_i, V_j, Z_j)) \\ &= \int P(R_{ijk} | U_i, V_j, D_k, Z_j) P(D_k) dD_k \\ &= \int f_{softmax}(U_i^T V_j + D_k W_d + Z_j W_z) P(D_k) dD_k \\ &= \int \frac{\exp(U_i^T V_j + Z_j W_z) \exp(D_k W_d)}{\sum_j \exp(U_i^T V_j + Z_j W_z) \exp(D_k W_d)} P(D_k) dD_k \\ &= \int \frac{\exp(U_i^T V_j + Z_j W_z)}{\sum_j \exp(U_i^T V_j + Z_j W_z)} P(D_k) dD_k \\ &= f_{softmax}(U_i^T V_j + Z_j W_z). \end{aligned} \quad (8)$$

Note that for $f_{softmax}(\cdot)$ we only consider items in the same domain.

3.4 Fine-tuning

Once the interaction data in the target domain is available, we can then fine-tune all the parameters of PREREC end-to-end in this new domain. During fine-tuning, we optimize Eqn. 7 on the target domain data and re-estimate all latent variables along with corresponding encoding functions as stated in Eqn. 5 and Eqn. 6.

For inference after fine-tuning, with the re-estimated latent vectors ($\widehat{D}_k, \widehat{Z}_j, \widehat{V}_j, \widehat{U}_i$) and all learnable parameters inside $f_{softmax}(\cdot)$, i.e., \widehat{W}_d and \widehat{W}_z , we then perform causal inference by intervening on the domain bias (D_k), popularity bias (Z_j), item properties (V_j), and user interests (U_i):

$$P(R_{ijk} | do(U_i, V_j, D_k, Z_j)) = f_{softmax}(\widehat{U}_i^T \widehat{V}_j + \widehat{D}_k \widehat{W}_d + \widehat{Z}_j \widehat{W}_z).$$

4 EXPERIMENTS

In this section, we first introduce the experiment set up, and then present the results and analysis with major goals to address the following questions:

- Q1** When training on multiple domains and testing on a new domain, will PREREC bring benefits to the new domain? What's the *zero-shot* performance of PREREC compared with baselines?
- Q2** How effective is the proposed casual debiasing mechanism to alleviate in- and cross-domain biases? To what extent does it improve the performance?
- Q3** Given the pre-trained model, how does PREREC perform compared to baselines if *full fine-tuning* (more details in Sec. 4.5) in the new domain is allowed?

Q4 Given the pre-trained model, if fine-tuning in the new domain is allowed, how does the *number of fine-tuning samples* affect final performance?

4.1 Dataset Processing and Statistics

We consider two datasets: (1) **XMarket** dataset³ [2], which is a large-scale real-world dataset covering 18 local markets (countries) on 16 different product categories; and (2) **Online Retail**⁴, which contains data collected from a UK-based online retail platform.

To ensure a sufficient number of non-overlapping users/items to evaluate zero-shot and fine-tuning performance rigorously while guaranteeing enough user-item interactions for training and evaluation, we set “India”, “Spain”, “Canada” in the XMarket dataset as the **pre-trained datasets**, set “Australia”, “Mexico”, “Germany”, “Japan” in the XMarket dataset as the **cross-market datasets**, and choose the online retail dataset as the **cross-platform dataset**. For experiments, we pre-train PREREC on the pre-trained datasets and evaluate it on both cross-market datasets and cross-platform datasets.

Domain and Dataset Split. For each domain/dataset, we randomly split users into training/validation/test sets with the ratio 4:3:3. The test splits of target domains are used for evaluation under both the zero-shot setting and the fine-tuning setting. Furthermore, to evaluate the zero-shot performance rigorously, we further filter out the overlapping users and items from the test splits; we name the remaining split **unseen** test set. See the supplement for the dataset statistics of each target domain.

4.2 Evaluated Methods

We compared our proposed PREREC with state-of-the-art methods, including **Random** (recommend by random selection), **POP** (recommend by popularity), **SBERT** [34], **GRU4Rec** [13], **SAS-Rec** [18], **ZESRec** [8], **UniSRec** [14], and **PreRec_n**. Please see the Appendix B for implementation details.

4.3 Evaluation Metrics

We use Recall@K% and our proposed r-NDCG@K% to evaluate model performance. For pre-training on multiple source domains, it is not ideal to evaluate performance and apply early stopping based on averaged Recall@K or NDCG@K. This is because different domains have different numbers of items, and domains with fewer items tend to have larger NDCG@K or Recall@K, consequently dominating the evaluation. Therefore, we use Recall@K% and our proposed r-NDCG@K% (‘r’ stands for ‘revised’) to do model evaluation and early-stop. Denoting the number of items in a domain as N , with the next interacted item as the single target, r-NDCG@K% is defined as follows:

$$\text{r-DCG} = 1/\log\left(a + \frac{b*(\text{rank}-1)}{N}\right), \quad (9)$$

$$\text{r-iDCG} = 1/\log\left(a + \frac{b*(1-1)}{N}\right) = 1/\log(a), \quad (10)$$

$$\text{r-NDCG} = \frac{\text{r-DCG}}{\text{r-iDCG}} = \frac{\log(a)}{\log\left(a + \frac{b*(\text{rank}-1)}{N}\right)}. \quad (11)$$

Here a and b are hyperparameters, and N is the number of items in a given domain. With this approach, the value of r-NDCG@K% is normalized by the number of items N in that domain. As a result, the value will not tend to be high when N is low. This a situation where the r-NDCG@K% of domain with small N dominates r-NDCG@K% of domain with large N .

Note that r-NDCG becomes the commonly used NDCG if $a = 2$ and $b = N$. In our experiments, we set $a = 2$ and $b = 15000$. Meanwhile, we use @K% instead of @K, with @K% indicating that only the top K% recommended items will contribute to the Recall@K% and r-NDCG@K%. Empirically, we set K%=0.04% for the cross-market scenario and set K%=0.5% for the cross-platform scenario. We use the average r-NDCG@K% on the validation set to early-stop training; we use both r-NDCG@K% and Recall@K% on the test set to evaluate a model.

4.4 Zero-shot Experiments

For the zero-shot experiment, we pre-train PREREC on three source domains (“India”, “Spain”, and “Canada” in the XMarket dataset) and evaluate its performance against a variety of baselines on five target domains (“Australia”, “Mexico”, “Germany”, “Japan” in the XMarket dataset and the Online Retail dataset) respectively.

Evaluation Setup. During evaluation, we simulate an online environment where models *can only access the target domain data during inference*, and all the models are evaluated on the test set of the target domain. Besides measuring performance on the whole test set, we also separately measure performance on the *unseen* test set (see Sec. 4.1). We utilize Recall@K% and r-NDCG@K% as evaluation metrics since they are less sensitive to differences in task difficulty across different domains. This is to ensure that performance for different domains is directly comparable, which is reflected by the consistent performance of the Random model across four target domains within the XMarket dataset (Recall@K% and r-NDCG@K% are constantly 0.004 and 0.002, respectively). The baselines can be broadly categorized into (1) **non-learnable zero-shot models**: Random, POP, SBERT, and (2) **learnable zero-shot models**: ZESRec, UnisRec, PreRec_n. Table 1 shows the zero-shot performance of different models.

Zero-shot Results (Q1). For the baselines, among the non-learnable zero-shot models, SBERT significantly outperforms both Random and POP since it leverages additional information-rich item descriptions. Importantly, SBERT achieves comparable performance with the learnable zero-shot models, which shows that without proper pre-training strategy to filter out the noise in the training data, the pre-training stage may not bring large performance gain due to overfitting on source domains. Among the learnable zero-shot models, the UnisRec slightly outperforms the ZESRec because it includes additional contrastive training objectives and uses a self-attention-based sequential model instead of GRU in ZESRec. Our proposed PREREC significantly outperforms all baselines by a large margin on almost all cases under both cross-market and cross-platform scenarios thanks to the causal debiasing mechanism which extracts the generic knowledge across source domains while reducing biases in data.

³<https://xmrec.github.io/>

⁴<https://www.kaggle.com/datasets/carrie1/ecommerce-data>

Scenario	Dataset	Metric	Test Type	Random	POP	SBERT	ZESRec	UniSRec	PreRec _n	PREREC
Cross -Market	Australia	Recall@K%	all	0.0004	0.0450	0.0552	0.0431	0.0404	<u>0.0583</u>	0.1036
			unseen	0.0004	0.0127	0.0619	0.0472	0.0452	<u>0.0655</u>	0.0756
		r-NDCG@K%	all	0.0002	0.0261	0.0380	0.0287	0.0281	<u>0.0396</u>	0.0656
			unseen	0.0002	0.0063	0.0420	0.0323	0.0316	<u>0.0438</u>	0.0476
	Mexico	Recall@K%	all	0.0004	0.0695	0.1509	0.1397	0.1521	<u>0.1645</u>	0.2316
			unseen	0.0004	0.0248	0.1403	0.1157	0.1386	0.1487	<u>0.1469</u>
		r-NDCG@K%	all	0.0002	0.0388	0.1095	0.0961	0.1134	<u>0.1178</u>	0.1557
			unseen	0.0002	0.0111	0.1009	0.0748	0.1000	0.1027	<u>0.1017</u>
	Germany	Recall@K%	all	0.0004	0.1514	0.2737	0.2639	0.2703	<u>0.2827</u>	0.3526
			unseen	0.0004	0.1016	<u>0.2528</u>	0.2295	0.2430	0.2506	0.2750
		r-NDCG@K%	all	0.0002	0.0936	<u>0.2103</u>	0.1886	0.2100	<u>0.2103</u>	0.2576
			unseen	0.0002	0.0559	0.1910	0.1703	0.1852	<u>0.1923</u>	0.1954
Japan	Recall@K%	all	0.0004	0.0657	<u>0.2926</u>	0.2259	0.2579	0.2817	0.3083	
		unseen	0.0004	0.0556	<u>0.2737</u>	0.2152	0.2662	0.2730	0.2741	
	r-NDCG@K%	all	0.0002	0.0376	<u>0.2016</u>	0.1514	0.1776	0.1929	0.2033	
		unseen	0.0002	0.0325	0.2061	0.1545	0.1967	0.1949	<u>0.1980</u>	
Cross -Platform	Online Retail	Recall@K%	all/unseen	0.0050	0.0672	0.1383	0.0564	0.1147	<u>0.1397</u>	0.1794
		r-NDCG@K%	all/unseen	0.0025	0.0362	<u>0.0848</u>	0.0316	0.0698	0.0841	0.1065

Table 1: Zero-shot performance comparison of different methods. K%= 0.04% for Cross-Market and K%= 0.5% for Cross-Platform. We mark the best results with bold face and the second best results with underline.

Scenario	Dataset	Metric	Random	POP	SBERT	GRU4Rec*	SASRec*	UniSRec	PREREC
Cross -Market	Australia	Recall@K%	0.0004	0.0450	0.0552	0.0546	<u>0.0735</u>	0.0600	0.1130
		r-NDCG@K%	0.0002	0.0261	0.0380	0.0349	<u>0.0495</u>	0.0416	0.0715
	Mexico	Recall@K%	0.0004	0.0695	0.1509	0.2315	0.2475	<u>0.2478</u>	0.2646
		r-NDCG@K%	0.0002	0.0388	0.1095	0.1683	0.1813	0.1832	<u>0.1825</u>
	Germany	Recall@K%	0.0004	0.1514	0.2737	0.3344	<u>0.3428</u>	0.3235	0.3935
		r-NDCG@K%	0.0002	0.0936	0.2103	0.2653	<u>0.2726</u>	0.2697	0.2964
	Japan	Recall@K%	0.0004	0.0657	0.2926	0.3428	<u>0.3645</u>	0.3587	0.3813
		r-NDCG@K%	0.0002	0.0376	0.2016	0.2475	<u>0.2554</u>	0.2549	0.2601
Cross -Platform	Online Retail	Recall@K%	0.0050	0.0672	0.1383	0.2524	<u>0.2728</u>	0.2513	0.2992
		r-NDCG@K%	0.0025	0.0362	0.0848	0.1545	<u>0.1665</u>	0.1589	0.1769

Table 2: Full fine-tuning results of different models. K%= 0.04% for Cross-Market and K%= 0.5% for Cross-Platform. The notation * indicates the model is trained from scratch. We mark the best results with bold face and the second best results with underline.

The improvement on unseen items is relatively lower than the overall improvement. We discovered that unseen items are usually unpopular items in target domains; this is also why POP’s performance is worse on “unseen” items than in “all” items.

For the zero-shot results of PREREC among cross-market datasets, we observed the largest improvement on Australia, moderate improvement on Germany and Mexico, with the least improvement on Japan. Such differences in improvements are potentially due to the differences in similarity between source domains and target domains: users in Australia share more common interests with users in Canada, which is the largest source domain; German speaks the same language as Canadian and Indian, and so does Mexican and Spanish; none of the source domains are Japanese-speaking countries, and Japanese users’ interests tend to be different from those in the source domains due to differences in culture.

Ablation Study on Causal Debiasing Mechanism (Q2). To verify the effectiveness of the causal debiasing mechanism, we

performed ablation studies on all four target domains. We implemented PreRec_n, a simplified version of PREREC, which ignores the cross-domain and in-domain bias terms during pre-training. Table 1 shows that in almost all cases PREREC outperforms PreRec_n.

4.5 Fine-tuning Experiments

For the fine-tuning experiments, we test under two settings: (1) *full fine-tuning* where we train or fine-tune models on all available training data in the target domain, (i.e., the training split of the target domain), and (2) *incremental fine-tuning* where we train or fine-tune models on a set of incrementally larger training data in the target domain. All the models are evaluated on the test set of the target domain. There are three sets of baselines: (1) **non-learnable zero-shot model**: Random, POP, SBERT; (2) **In-domain model**: GRU4Rec, SASRec; (3) **Pre-trained model**: UniSRec.

Full Fine-tuning (Q3). Table 2 shows the full fine-tuning performance of different methods. All in-domain models (GRU4Rec and

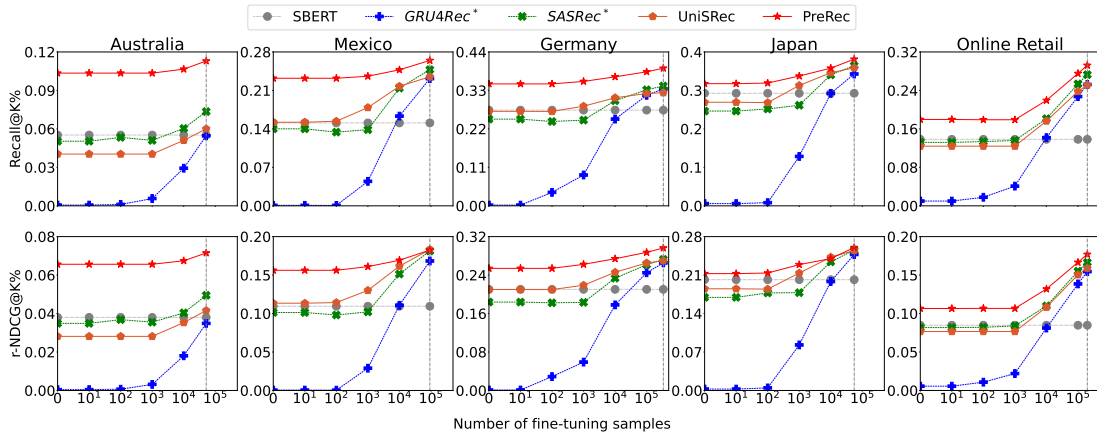


Figure 3: Incremental fine-tuning results of different models. $K\%= 0.04\%$ for Cross-Market and $K\%= 0.5\%$ for Cross-Platform. The solid line indicates the model is pre-trained while the dashed line indicates the model is trained from scratch. The last point for each line corresponds to the full fine-tuning with all available training data in the target domain.

SASRec) achieve better performance than SBERT, which is the best non-learnable zero-shot model. The pre-trained model UniSRec achieves performance similar to the in-domain models, showing that UniSRec brings no additional gain from pre-training on other domains if there is sufficient data in target domains. PREREC outperforms all baselines on almost all cases under both cross-market and cross-platform scenarios by a large margin after fine-tuning, demonstrating that: (1) PREREC distills generic knowledge in the pre-training stage, and it still has complementary value even after fine-tuning on target domain data; (2) PREREC is capable of rapidly adapting to the new domain without forgetting the distilled knowledge during pre-training (i.e., robust to catastrophic forgetting).

Incremental Fine-tuning (Q4). Fig 3 shows the results for incremental fine-tuning. Among all target domains, PREREC outperforms all baselines, with the performance steadily improving as the number of target domain training samples increases, showing PREREC is progressively adapting to the target domain; the performance gap between PREREC and the baselines is gradually shrinking, suggesting diminishing complementary value of the distilled knowledge during pre-training. It is worth noting that the performance gap between PREREC and baselines is still prominent even when there are as many as 10^4 target domain training samples, equivalent to several weeks/months of effort on collecting data.

5 RELATED WORK

In-domain Recommendation. There is a rich literature on in-domain recommendation, i.e., training and testing the recommender systems on the same domain. Collaborative filtering methods such as PMF [31] and BPR [35] are first proposed to approach this problem. Later the deep learning methods such as GRU [5], Transformer [37] and Graph Neural Network [47] were proposed and demonstrated superiority on a variety of tasks, works such as GRU4Rec [13], SAS4Rec [18] and KGAT [43] adopt these latest neural network architectures and achieve great success in the recommendation regime. These methods assume only rating/interaction data is available without any content information (e.g., text description for items). Such an assumption precluded their direct application to pretraining-based recommenders. A new line of deep

learning recommenders, pioneered by collaborative deep learning (CDL) [40] and its variants [23, 39], seamlessly incorporate content information into deep recommenders, thereby opening up the possibility of pretraining-based recommenders and significantly alleviating cold-start problems in recommender systems.

Cross-domain Recommendation. Beyond in-domain recommendation, there are other works that build on domain adaptation methods [11, 27, 36, 38, 48, 49] from the machine learning community to achieve cross-domain recommendation. Specifically, they utilize data from source domains to boost recommendation performance in target domains with either common users or common items [1, 12, 20, 24, 26, 46, 50, 59]. Another line of work, usually referred to as Dual-Target Cross-Domain Recommendation, tries to improve performance of both source and target domains [16, 21, 25, 55] (a more comprehensive survey of cross-domain recommender systems can be found in [58]). Note that these aforementioned cross-domain recommendation methods are fundamentally different from pre-trained recommendation.

Bias in Recommendation. While recommender systems are analyzed through in-domain and cross-domain perspectives, there is an emphasis on understanding them through the lens of bias [4]. Within a single domain, the topic of item popularity debiasing has garnered significant attention. Studies [30, 54, 56] have demonstrated that addressing item popularity bias can enhance recommendation quality. In the broader context of cross-domain bias mitigation, [22] has adapted traditional Inverse-Propensity-Score (IPS) to fit cross-domain contexts, while [9] employs invariant representations to reduce the impact of spurious correlations. However, the method to effectively leverage item popularity knowledge from pre-training domains to new domains remains an open question.

Causality in Recommendation. Closely related to debiasing are studies on causality-inspired recommender systems. [44, 45] extract invariant user-item embeddings across varying scenarios, aiming to minimize the influence of potentially misleading features. Both [53] and [57] integrate the influence of popularity into their

causal analyses within a specific domain. Nonetheless, the application of causality on a broader scale, specifically within pre-trained recommender systems, remains largely unexplored.

6 CONCLUSION

In this paper, we identify two types of bias: (1) in-domain bias: we consider the item popularity bias taking effects within each domain; and (2) cross-domain bias: we consider variability introduced by the unique domain properties. We propose PREREC equipped with a novel causal debiasing mechanism to deal with both bias terms. We extensively conduct experiments under both cross-market and cross-platform scenarios and demonstrate the effectiveness of our model. Future work includes identifying other confounders to incorporate into our causal debiasing framework, experimenting with better PLMs such as OPT [51], exploring other modalities, etc.

REFERENCES

- [1] Ye Bi, Liqiang Song, Mengqiu Yao, Zhenyu Wu, Jianming Wang, and Jing Xiao. 2020. A Heterogeneous Information Network based Cross Domain Insurance Recommendation System for Cold Start Users. In *SIGIR*. 2211–2220.
- [2] Hamed Bonab, Mohammad Aliannejadi, Ali Vardasbi, Evangelos Kanoulas, and James Allan. 2021. Cross-market product recommendation. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 110–119.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [4] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems* 41, 3 (2023), 1–39.
- [5] Junyoung Chung, cCaglar Gülcebre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *CoRR* abs/1412.3555 (2014). <http://arxiv.org/abs/1412.3555>
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Hao Ding, Branislav Kveton, Yifei Ma, Youngsuk Park, Venkataramana Kini, Yupeng Gu, Ravi Divvela, Fei Wang, Anoop Deoras, and Hao Wang. 2023. Trending Now: Modeling Trend Recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 294–305.
- [8] Hao Ding, Yifei Ma, Anoop Deoras, Yuyang Wang, and Hao Wang. 2021. Zero-shot recommender systems. *arXiv preprint arXiv:2105.08318* (2021).
- [9] Xiaoyu Du, Zike Wu, Fuli Feng, Xiangnan He, and Jinhui Tang. 2022. Invariant Representation Learning for Multimedia Recommendation. In *MM '22: The 30th ACM International Conference on Multimedia*. 619–628.
- [10] Ziwei Fan, Hao Ding, Anoop Deoras, and Trong Nghia Hoang. 2023. Personalized federated domain adaptation for item-to-item recommendation. In *Proceedings of the UAI 2023*. <https://www.amazon.science/publications/personalized-federated-domain-adaptation-for-item-to-item-recommendation>
- [11] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, Franois Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research* 17, 1 (2016), 2096–2030.
- [12] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Content-aware Neural Hashing for Cold-start Recommendation. In *SIGIR*. 971–980.
- [13] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [14] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards Universal Sequence Representation Learning for Recommender Systems. *CoRR* abs/2206.05941 (2022). <https://doi.org/10.48550/arXiv.2206.05941>
- [15] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards Universal Sequence Representation Learning for Recommender Systems. *arXiv preprint arXiv:2206.05941* (2022).
- [16] Guangneng Hu, Yu Zhang, and Qiang Yang. 2018. CoNet: Collaborative Cross Networks for Cross-Domain Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM*. 667–676.
- [17] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately interpreting clickthrough data as implicit feedback. In *Acm Sigir Forum*, Vol. 51. Acm New York, NY, USA, 4–11.
- [18] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [19] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [20] Jingjing Li, Mengmeng Jing, Ke Lu, Lei Zhu, Yang Yang, and Zi Huang. 2019. From zero-shot learning to cold-start recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4189–4196.
- [21] Pan Li and Alexander Tuzhilin. 2020. DDTCDR: Deep Dual Transfer Cross Domain Recommendation. In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3–7, 2020*, James Caverlee, Xia (Ben) Hu, Mounia Lalmas, and Wei Wang (Eds.). ACM, 331–339. <https://doi.org/10.1145/3336191.3371793>
- [22] Siqing Li, Liuyi Yao, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Tonglei Guo, Bolin Ding, and Ji-Rong Wen. 2021. Debiasing Learning based Cross-domain Recommendation. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3190–3199.
- [23] Xiaopeng Li and James She. 2017. Collaborative variational autoencoder for recommender systems. In *KDD*. 305–314.
- [24] Tingting Liang, Congying Xia, Yuyu Yin, and Philip S. Yu. 2020. Joint Training Capsule Network for Cold Start Recommendation. In *SIGIR*. 1769–1772.
- [25] Meng Liu, Jianjun Li, Guohui Li, and Peng Pan. 2020. Cross Domain Recommendation via Bi-directional Transfer Graph Collaborative Filtering Networks. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management*. 885–894.
- [26] Siwei Liu, Iadh Ounis, Craig Macdonald, and Zaiqiao Meng. 2020. A Heterogeneous Graph Neural Model for Cold-Start Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2029–2032.
- [27] Tianyi Liu, Zihao Xu, Hao He, Guangyuan Hao, Guang-He Lee, and Hao Wang. 2023. Taxonomy-Structured Domain Adaptation. In *ICML*.
- [28] Yiming Liu, Xuezhi Cao, and Yong Yu. 2016. Are you influenced by others when rating? Improve rating prediction by conformity modeling. In *Proceedings of the 10th ACM conference on recommender systems*. 269–272.
- [29] Yifei Ma, Balakrishnan Narayanaswamy, Haibin Lin, and Hao Ding. 2020. Temporal-contextual recommendation in real-time. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2291–2299.
- [30] Yifei Ma, Balakrishnan (Murali) Narayanaswamy, Haibin Lin, and Hao Ding. 2020. Temporal-Contextual Recommendation in Real-Time. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23–27, 2020*, Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash (Eds.). ACM, 2291–2299. <https://doi.org/10.1145/3394486.3403278>
- [31] Andriy Mnih and Russ R Salakhutdinov. 2007. Probabilistic matrix factorization. *Advances in neural information processing systems* 20 (2007).
- [32] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [33] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. <http://arxiv.org/abs/1908.10084>
- [34] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. <http://arxiv.org/abs/1908.10084>
- [35] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [36] Haizhou Shi and Hao Wang. 2023. A Unified Approach to Domain Incremental Learning with Memory: Theory and Algorithm.. In *NeurIPS*.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [38] Hao Wang, Hao He, and Dina Katabi. 2020. Continuously Indexed Domain Adaptation. In *ICML*.
- [39] Hao Wang, Xingjian Shi, and Dit-Yan Yeung. 2016. Collaborative recurrent autoencoder: Recommend while learning to fill in the blanks. In *NIPS*. 415–423.
- [40] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 1235–1244.
- [41] Hao Wang and Dit-Yan Yeung. 2016. Towards Bayesian deep learning: A framework and some existing methods. *TDKE* 28, 12 (2016), 3395–3408.
- [42] Hao Wang and Dit-Yan Yeung. 2020. A Survey on Bayesian Deep Learning. *CSUR* 53, 5 (2020), 1–37.

- [43] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 950–958.
- [44] Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. 2020. Causal inference for recommender systems. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 426–431.
- [45] Zimu Wang, Yue He, Jiashuo Liu, Wenchao Zou, Philip S. Yu, and Peng Cui. 2022. Invariant Preference Learning for General Debiasing in Recommendation. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1969–1978.
- [46] Le Wu, Yonghui Yang, Lei Chen, Defu Lian, Richang Hong, and Meng Wang. 2020. Learning to Transfer Graph Embeddings for Inductive Graph based Recommendation. In *SIGIR*. 1211–1220.
- [47] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 346–353.
- [48] Zihao Xu, Guangyuan Hao, Hao He, and Hao Wang. 2023. Domain Indexing Variational Bayes: Interpretable Domain Index for Domain Adaptation. In *ICLR*.
- [49] Zihao Xu, Guang-He Lee, Yuyang Wang, Hao Wang, et al. 2022. Graph-Relational Domain Adaptation. In *ICLR*.
- [50] Fajie Yuan, Xiangnan He, Alexandros Karatzoglou, and Liguang Zhang. 2020. Parameter-Efficient Transfer from Sequential Behaviors for User Modeling and Recommendation. In *SIGIR*. 1469–1478.
- [51] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068* (2022).
- [52] Yuhui Zhang, Hao Ding, Zeren Shui, Yifei Ma, James Zou, Anoop Deoras, and Hao Wang. 2021. Language Models as Recommender Systems: Evaluations and Limitations. In *I (Still) Can't Believe It's Not Better! NeurIPS 2021 Workshop*. <https://openreview.net/forum?id=hFx3FY7-m9b>
- [53] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal intervention for leveraging popularity bias in recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 11–20.
- [54] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.). ACM, 11–20. <https://doi.org/10.1145/3404835.3462875>
- [55] Cheng Zhao, Chenliang Li, and Cong Fu. 2019. Cross-Domain Recommendation via Preference Propagation GraphNet. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM*. 2165–2168.
- [56] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Yong Li, and Depeng Jin. 2020. Disentangling User Interest and Popularity Bias for Recommendation with Causal Embedding. *CoRR abs/2006.11011* (2020). [arXiv:2006.11011](https://arxiv.org/abs/2006.11011) <https://arxiv.org/abs/2006.11011>
- [57] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Yong Li, and Depeng Jin. 2021. Disentangling user interest and conformity for recommendation with causal embedding. In *Proceedings of the Web Conference 2021*. 2980–2991.
- [58] Feng Zhu, Yan Wang, Chaochao Chen, Jun Zhou, Longfei Li, and Guanfang Liu. 2021. Cross-Domain Recommendation: Challenges, Progress, and Prospects. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, Zhi-Hua Zhou (Ed.). ijcai.org, 4721–4728. <https://doi.org/10.24963/ijcai.2021/639>
- [59] Ziwei Zhu, Shahin Sefati, Parsa Saadatpanah, and James Caverlee. 2020. Recommendation for New Users and New Items via Randomized Training and Mixture-of-Experts Transformation. In *SIGIR*. 1121–1130.

Pre-trained Recommender Systems: A Causal Debiasing Perspective (Supplementary Material)

A NEURAL NETWORK ARCHITECTURE

Figure 4 shows the neural network architecture of our PREREC. The item overall property \mathbf{V}_{jk} is a summation of domain property \mathbf{D}_k , item popularity property \mathbf{Z}_{jk} , and item embedding extracted from item description via a pre-trained frozen multilingual language model and a linear layer. A sequential model will aggregate item embeddings of items in user history to generate user embedding U_{ik} via multi-head attention. Then user/item similarity is the inner product between U_{ik} and \mathbf{V}_{jk} . This similarity will be further calibrated by popularity score and domain bias to form the use/item preference \mathbf{R}_{ijk} .

B IMPLEMENTATION DETAILS

Here we note down the details for model implementations and the training process.

We consider the following methods during evaluation:

- **Random**: Recommend items by random selection from the whole item catalogue without replacement.
- **POP**: Recommend items based on the items' popularity in the last stage (i.e., the last 15 days in our experiments).
- **SBERT**: Apply the pre-trained SBERT [34] on the item's description and generate the textual embedding as the item embedding, use the item embedding of the last interacted item in the user sequence as the user embedding, and recommend the next item to the user based on the inner product between the user embedding and the item embedding.
- **GRU4Rec** [13]: Use GRU to model the user interaction sequence for session-based recommendation. For fair comparison, we use the SBERT embedding generated from item's description as the input item representation.
- **SASRec** [18]: A self-attention based sequential model for session-based recommendation. For fair comparison, we use the SBERT embedding generated from item's description as the input item representation.
- **ZESRec** [8]: A flexible sequential framework that can be pre-trained on a source domain and directly applied to a target domain. It generates the universal item embedding via pre-trained language models (in our experiments we adopt SBERT) and generates universal user embeddings by aggregating universal item embeddings in the user sequence. In the experiments, we employ the GRU variant of ZESRec.
- **UniSRec** [14]: A follow-up work of ZESRec which further enables pre-training on multiple source domains. It adopts a self-attention based sequential model, equipping it with a MoE-enhanced Adaptor and additional contrastive training objectives, to assist with domain fusion and improve performance. As [14] claims, UniSRec is a state-of-the-art cross-domain model.
- **PreRec_n**: A simplified version of PREREC without the causal debiasing mechanism; in other words it does not take the cross-domain bias and the in-domain bias into the consideration during pre-training.

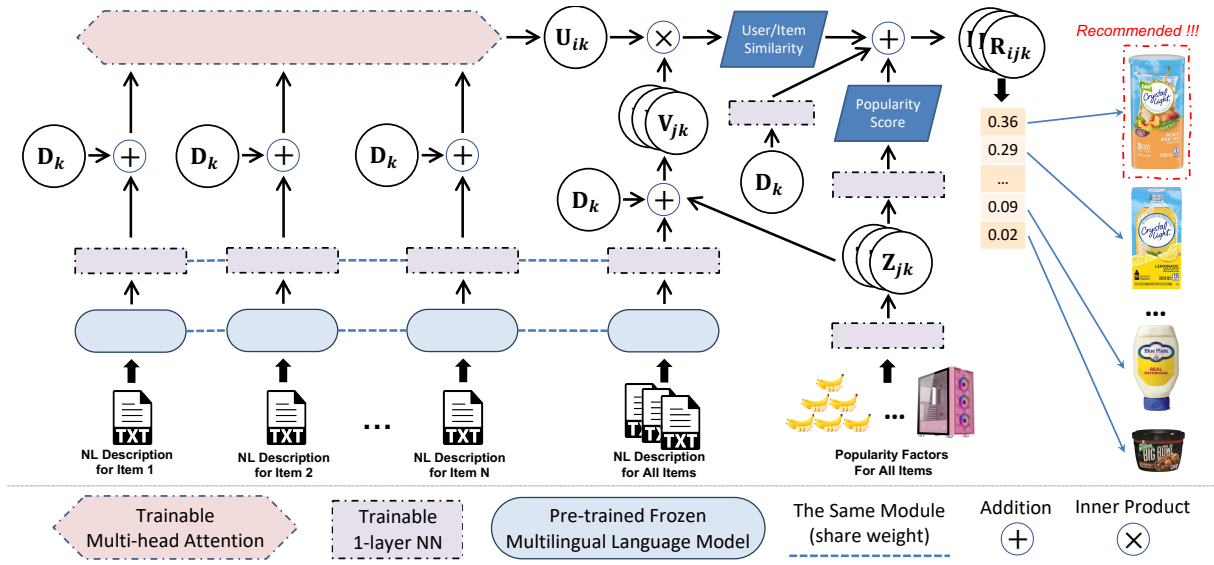


Figure 4: The neural network implementation for our PREREC. Analogous to Fig. 2 in the main paper, f_{BERT} is implemented as a multilingual language model, f_{seq} is implemented as a transformer decoder, and f_{pop} is implemented as a linear layer with activation.

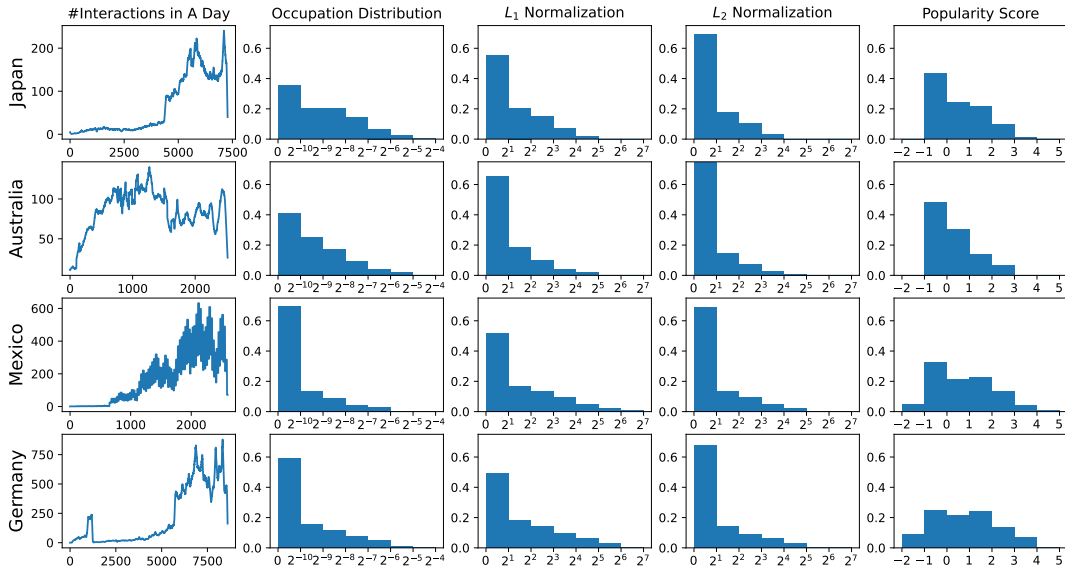


Figure 5: The illustration of how popularity score works. The first column shows different domains have different traffic volumes and the second column shows occupation is not comparable among different domains since different domains have different numbers of available items. The third and fourth columns show the distribution of the first and the second dimension of the popularity factor F in each domain, which is a normalized frequency proposed in this paper, and shown to be more comparable among different domains. After perturbing, the popularity factor F is further mapped to the popularity score. As shown in the last column, the popularity score is comparable among zero-shot domains, indicating the generalizability of our proposed method for popularity.

Models. All methods except Random and POP use SentenceBERT [34] to extract item BERT embeddings. For a fair comparison, all methods use 256-dimensional item embeddings and user embeddings; for sequential models based on GRU, the number of layers

of GRU is set to 2; for self-attention based sequential models, the number of multi-head attention layers is set to 2.

Training. We use Adam [19] for both pre-training and fine-tuning with the learning rate of 0.0003. L2 regularization is implemented with weight decay in PyTorch’s Adam. We use the average

r-NDCG@K% on the validation set to early-stop pre-training, after which the pre-trained model is used for fine-tuning on each target domain. See the supplement for more details.

C THE DESIGN RATIONALE OF COMPUTING POPULARITY PROPERTIES

To calculate the popularity score in a domain k , assuming a set of items J_k , and the number of interactions of an item j in time slot T is c_j^T . Naively, one could compute popularity factors F_j via the following two options:

- **Method A:** Directly using $\log c_j^T$ or c_j^T as the popularity factors F_j for item j at time interval T .
- **Method B:** Computing the percentage of the traffic volume, i.e., using $c_j^T / \sum_{i \in J_k} c_i^T$ or $(c_j^T / \sum_{i \in J_k} c_i^T)^p$ as the popularity factors F_j for item j at time interval t , where p is a pre-defined scalar.

However, these two methods are not transferable across multiple domains:

- For **Method A**, assuming domain m has more interactions than domain n and both domains have the same number of items. In this case, even though item j in domain m has more interactions than item i in domain n , it does not mean that item j in domain m is more popular than item i in domain n .
- For **Method B**, assuming domain m has more items than domain n . In this case, even if item j in domain m has the same percentage of traffic volume (e.g. 5%) as item i in domain n , item j may still be more popular in domain m than item i in domain n .

To address the abovementioned issue, we design a method inspired by the half Gaussian distribution. Assuming in any domain, the number of interactions of each item follows a half Gaussian distribution. The popularity factor F_j of item j can be measured as:

$$F_j = c_j^T / \sqrt{\frac{\sum_{i \in J_k} (c_i^T)^2}{|J_k|}}, \quad (12)$$

where the popularity factor F_j is transferable from domain to domain.

We check the distribution over the number of interactions of items for each real-world dataset and discover that they may not precisely follow the half Gaussian distribution. Therefore, we propose to consider different orders of norms for the denominator in Eqn. 12 and compute a set of popularity factors as F_j :

$$F_j = \left[\frac{c_j^T}{s_1^T}, \frac{c_j^T}{s_2^T}, \dots, \frac{c_j^T}{s_w^T} \right],$$

where s_w^T is a normalization term calculated as:

$$s_w^T = \left(\sum_{j \in J_k} (c_j^T)^w / |J_k| \right)^{\frac{1}{w}}$$

with varied k as the popularity property, and use a trainable neural network to get the popularity score. The final w ranges from 1 to 4, i.e., the popularity property factors F_j is a 4-dimension vector.

Fig. 5 illustrates how the popularity score works on four zero-shot domains.

D THE CHOICE OF K% IN R-NDCG@K% AND RECALL@K%

We choose $K\% = 0.04\%$ due to under the number of samples N from Table 3 we have $10 \leq N \times K\% \leq 50$. This provides the results between NDCG@10 and NDCG@50, which aligns with common evaluation setting including NDCG@10, NDCG@20, or NDCG@50.

E ADDITIONAL TRAINING DETAILS

We use one Tesla V100 GPU for training the model. The pre-training phase takes around 25 minutes per epoch, and the maximum number of training epochs is 30.

For contrastive learning, UniSRec uses in-batch negative. Thus, the negative items are from all domains. PREREC adopts a more complex implementation, sampling negative items randomly and uniformly from the same domain as the positive item.

For hyperparameter tuning, we consider candidates as shown in Table 4.

Data	Pre-train			Zero-shot/Fine-tune				
	India	Spain	Canada	Australia	Mexico	Germany	Japan	Online Retail
Laguage	English	Spain	English	English	Spain	English	Japanese	English
#Items	45893	39675	99376	42094	43095	70527	22591	4223
#Users	507581	400883	992366	86975	249229	997555	277570	24446
#Inters	748607	708103	1971956	213086	483660	1840912	465746	540455
#Unseen	/	/	/	14525	32477	60637	12831	516009

Table 3: Dataset statistics for different domains of XMarket and Online Retail. Language indicates the default language used on Online Retail and the Amazon website for that domain (country). #Inters is the number of interactions, #Unseen is the number of target-domain samples that are non-overlapping with source domains.

hyperparameter	candidate value
learning rate	0.00003, 0.0003 , 0.003, 0.03
batch size	32, 64, 128, 256 , 512
# of negtive samples	63, 127, 255 , 511
# of gru layers	1, 2 , 3, 4, 6, 8, 12
# of attention layers	1, 2 , 3, 4, 6, 8, 12
# of attention heads	1, 2 , 3, 4, 6, 8, 12
# of embedding dimensions for user/item	64, 128, 256 , 512, 1024
domain bias regularization	0.03, 0.3 , 3, 30, 300
item bias regularization	100

Table 4: The hyperparameter tuning setting. The selected ones are marked in bold face.