

# Reformulating CTR Prediction: Learning Invariant Feature Interactions for Recommendation

Yang Zhang\*  
University of Science and Technology  
of China  
Hefei, China  
zy2015@mail.ustc.edu.com

Tianhao Shi\*  
University of Science and Technology  
of China  
Hefei, China  
sth@mail.ustc.edu.com

Fuli Feng<sup>†</sup>  
University of Science and Technology  
of China  
Hefei, China  
fulifeng93@gmail.com

Wenjie Wang  
National University of Singapore  
Singapore  
wenjiewang96@gmail.com

Dingxian Wang  
Etsy Inc.  
United States  
dingxianwang@etsy.com

Xiangnan He<sup>†</sup>  
University of Science and Technology  
of China  
Hefei, China  
xiangnanhe@gmail.com

Yongdong Zhang  
University of Science and Technology  
of China  
Hefei, China  
zhyd73@ustc.edu.cn

## ABSTRACT

Click-Through Rate (CTR) prediction plays a core role in recommender systems, serving as the final-stage filter to rank items for a user. The key to addressing the CTR task is learning feature interactions that are useful for prediction, which is typically achieved by fitting historical click data with the Empirical Risk Minimization (ERM) paradigm. Representative methods include Factorization Machines and Deep Interest Network, which have achieved wide success in industrial applications. However, such a manner inevitably learns *unstable feature interactions*, *i.e.*, the ones that exhibit strong correlations in historical data but generalize poorly for future serving.

In this work, we reformulate the CTR task — instead of pursuing ERM on historical data, we split the historical data chronologically into several periods (*a.k.a.* environments), aiming to learn feature interactions that are **stable** across periods. Such feature interactions are supposed to generalize better to predict future behavior data. Nevertheless, a technical challenge is that existing invariant learning solutions like Invariant Risk Minimization are not applicable, since the click data entangles both environment-invariant and environment-specific correlations. To address this dilemma,

we propose *Disentangled Invariant Learning* (DIL) which disentangles feature embeddings to capture the two types of correlations separately. To improve the modeling efficiency, we further design LightDIL which performs the disentanglement at the higher level of the feature field. Extensive experiments demonstrate the effectiveness of DIL in learning stable feature interactions for CTR.

## CCS CONCEPTS

• Information systems → Recommender systems.

## KEYWORDS

Factorization Machine; Recommender System; Invariant Learning

### ACM Reference Format:

Yang Zhang, Tianhao Shi, Fuli Feng, Wenjie Wang, Dingxian Wang, Xiangnan He, and Yongdong Zhang. 2023. Reformulating CTR Prediction: Learning Invariant Feature Interactions for Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539618.3591755>

## 1 INTRODUCTION

Click-Through Rate (CTR) prediction is important to support the ranking stage of recommender systems [43]. Its key lies in learning the feature interactions that are useful for predicting user clicks [11, 43, 46]. Existing methods [4, 11, 14, 27, 32, 43, 45] typically achieve the target by fitting historical data collectively with the Empirical Risk Minimization (ERM) paradigm — minimizing the average empirical loss over historical click data [43]. The ERM paradigm has become the standard paradigm for training CTR models, leading to the classical solutions FM [32], NFM [14], and DIN [45] that have been intensively used in industrial applications [43].

We argue that the CTR model built with the ERM paradigm would inevitably learn *unstable feature interactions*, leading to poor generalization for future serving. Some feature interactions could

\*Equal contribution.

<sup>†</sup>Corresponding authors.


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](mailto:permissions@acm.org).

*SIGIR '23, July 23–27, 2023, Taipei, Taiwan*

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9408-6/23/07...\$15.00

<https://doi.org/10.1145/3539618.3591755>



**Figure 1: Examples of stable feature interaction (gender, category) and unstable feature interaction (stock, category) and (price, color). The arrow represents that the feature interaction affects the happening of click, and the weights reflect the influence strength. The dotted line represents the correlations between feature interactions and clicks, varying along the timeline.**

exhibit strong correlations in historical data but are useless or even harmful for future serving. Figure 1 illustrates two examples of unstable feature interactions: 1) the feature interaction (stock, category) highly affects user clicks during the COVID-19 lockdown, however, its effect diminishes after the lockdown; and 2) the feature interaction (price, color) exhibits high correlation with clicks due to some confounding effect<sup>1</sup>, which is however unstable and becomes weak after the occurrence of counterfeits. These unstable feature interactions will mislead the CTR model learned by ERM, e.g., 1) over-recommending emergent supplies after the lockdown; and 2) over-recommending the counterfeits with high prices and attractive colors. It is thus essential to distinguish unstable feature interactions in CTR prediction.

In this work, we propose to reformulate the CTR task instead of performing ERM on historical data. As unstable feature interactions could cause poor generalization, we set the target as learning the feature interactions that are stably useful for predicting user clicks. Towards the goal, we split the historical data chronologically into several periods, forming a set of environments, and learn feature interactions that are stable across these environments. These environments could be heterogeneous due to the temporal influence, e.g., some environments have more historical data during the COVID-19 lockdown. The feature interactions that are stably useful over these heterogeneous environments are more likely to remain useful in the future [2, 25], e.g., the interaction (gender, category) in Figure 1. As such, the CTR models accounting for the stability of feature interactions will generalize better for future prediction.

Among existing techniques, Invariant Learning [2] is well-known for capturing environment-invariant correlations across heterogeneous environments, which seems to be a promising choice for learning stable feature interactions in CTR prediction. However, it is not directly applicable, because it relies on the sufficiency prediction condition – the prediction target can be sufficiently predicted with only environment-invariant correlations – which is not satisfied in CTR prediction. Although unstable, some feature interactions truly affect the clicks in a specific environment, e.g., during the COVID-19 lockdown, users click items by considering the (stock, category) factor. In other words, the clicks entangle both environment-invariant and environment-specific correlations, meaning that the clicks cannot be sufficiently predicted with only environment-invariant correlations.

<sup>1</sup>E.g., the high quality of the product could synchronously lead to high item price, attractive item color, and strong click probability.

To overcome this challenge, we propose *Disentangled Invariant Learning* (DIL), which disentangles the environment-invariant and environment-specific correlations from clicks, and identifies stable feature interactions by avoiding the influence of environment-specific correlations. Specifically, DIL first equips a CTR model with environment-invariant feature embedding and a set of environment-specific feature embeddings. Then DIL optimizes the model with a new disentangled invariant learning objective, which contains: 1) environment-specific objectives to make environment-specific feature embeddings capture environment-specific correlations, and 2) a modified invariant learning objective with the sufficiency prediction condition satisfied to let environment-invariant feature embeddings capture environment-invariant correlations. Considering that embedding disentanglement significantly increases the model size and the cost of training, we further design LightDIL to perform the disentanglement at the higher level of the feature field to improve the modeling efficiency.

The main contributions are summarized as follows:

- **New Problem:** It is the first time that the CTR task is reformulated by learning stable feature interactions across periods, so as to achieve better generalization for future serving.
- **New Technique:** We propose to integrate the idea of representation disentanglement into invariant learning, making invariant learning feasible in difficult situations where the sufficiency prediction condition is not satisfied.
- **Experiments:** We conduct extensive experiments on both semi-synthetic and real-world datasets, verifying the effectiveness of our proposal.

## 2 PRELIMINARIES

This work studies the CTR task on chronologically collected data. Assuming the data is collected from  $T$  periods, which could be days, weeks, or self-defined time spans. Let  $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_T\}$  denote the data, where each  $\mathcal{D}_t$  represents the data collected at the  $t$ -th period. We denote each sample in  $\mathcal{D}$  as  $(x, y)$ , where  $y \in \{0, 1\}$  represents the click and  $x = [x_1, \dots, x_N]^T \in \{0, 1\}^N$  represents the features describing the user-item pair.  $x$  is a highly sparse multi-hot vector with a dimensionality of  $N$ . We assume the feature vector contains  $M$  fields (e.g., age and category) where each field is encoded by one-hot encoding or multi-hot encoding. The target is to learn a CTR model from  $\mathcal{D}$  to serve for future periods. In this work, we propose to perform stable feature interaction learning to enhance the generalization of the classic FM model. Next, we introduce the backbone FM model and the basics of invariant learning.

### 2.1 Factorization Machines

FM and its variants have achieved wide success in CTR, owing to the effectiveness of the inner product operation in explicitly modeling feature interactions [11, 14, 32]. Without losing generality, we take FM as the base model to study stable feature interaction modeling. Given an input sample with features  $x = [x_1, \dots, x_N]^T$ , FM generates the prediction by modeling all second-order feature interactions via inner product on pairwise feature embeddings:

$$\hat{y} = \sum_{i=1}^N \sum_{j>i}^N \langle v_i, v_j \rangle \cdot x_i x_j, \quad (1)$$

where  $\hat{y}$  denotes the prediction of CTR;  $v_i/v_j$  denotes the embedding for  $i$ -th/ $j$ -th feature, and  $\langle v_i, v_j \rangle$  refers to the inner product between

the two feature embeddings to reconstruct the interaction effect. The prediction accounts for the interaction between features  $x_i$  and  $x_j$  if both features occur, *i.e.*,  $x_i x_j = 1$ .

• **Field-level.** Since each feature belongs to one field, we can also model the interaction at the field level. Formally,

$$\hat{y} = \sum_{i=1}^M \sum_{j>i}^M \langle \mathbf{u}_i, \mathbf{u}_j \rangle, \quad (2)$$

where  $\mathbf{u}_i$  and  $\mathbf{u}_j$  denote the representation of the  $i$ -th and  $j$ -th fields. When all fields are one-hot encoding, field-level interaction modeling is equivalent to Equation (1). Let  $\mathcal{S}_i$  denote the feature indices belonging to the  $i$ -th field, we can generate the field representation  $\mathbf{u}_i$  based on the feature embeddings  $\{v_k\}_{k \in \mathcal{S}_i}$  as  $\mathbf{u}_i = \frac{1}{\sum_{k \in \mathcal{S}_i} x_k} \sum_{k \in \mathcal{S}_i} x_k \cdot v_k$ .

## 2.2 Invariant Learning

Invariant learning [36] is widely used to learn predictors for Out-of-Distribution (OOD) generalization, given the data collected from different environments. Existing invariant learning methods typically have the following assumption [2, 25, 41] (here, we reuse  $t$  to denote an environment since a period represents an environment):

**ASSUMPTION 1.** *There exist random variables  $\mathbf{r}$  that satisfy:*

- *Invariance condition: for any two environments  $t$  and  $t'$ ,  $P(y|\mathbf{r}, t) = P(y|\mathbf{r}, t')$  holds, where  $y$  denotes the prediction target, *e.g.*, the click.*
- *Sufficiency condition:  $y = f(\mathbf{r}) + \epsilon$ , where  $f$  represents a function and  $\epsilon$  is an independent noise.*

This assumption means that capturing the correlations between the predictor variables  $\mathbf{r}$  and the prediction target  $y$  can result in a desired model with good generalization to *unknown* distributions [18], because these correlations are invariant and sufficient for predicting the target across environments. As such, they are also denoted as *environment-invariant correlations*. Under this assumption, to make a model capture such correlations, one representative method is the variance-based regularizer [25, 41], which can be formulated as follows:

$$\min_{\phi} \frac{1}{T} \sum_{t=1}^T R^t + \lambda V_R, \quad (3)$$


where  $\phi$  denotes the model parameters,  $R^t$  represents the loss for the environment  $t$ ,  $T$  represents the number of environments,  $\lambda$  is a hyper-parameter to control the strength of  $V_R$ , and  $V_R$  denotes the variance of the loss [41] or the variance of the gradient of the loss over environments [25]. In this work, we adopt the variance of the loss with the consideration of computation efficiency, which is computed as follows:

$$V_R = \frac{1}{T-1} \sum_{t=1}^T \left( R^t - \frac{1}{T} \sum_{t'=1}^T R^{t'} \right)^2. \quad (4)$$

The key consideration of the method is that environment-invariant correlations should enable better performance across environments.

## 3 METHODOLOGY

Given the possible existence of unstable feature interaction, we first rethink the CTR prediction task to shed more light on the new problem setting. We then elaborate on the proposed disentangled invariant learning solution.



**Figure 2: Graphical model of the click generation process. The box represents  $T$  replicates of the part of the model in the box, indicating that the part varies over  $T$  periods.**

### 3.1 Rethinking the CTR Task

Figure 2 illustrates the generative process of click data with a graphical model. The box in the figure represents  $T$  replicates of the content in the box, indicating that the content could vary over  $T$  periods, *e.g.*,  $H^2 \rightarrow Y$  could vary over  $T$  periods. We divide the feature interactions into three major types  $H^1, H^2, H^3$ :

- $H^1$  denotes the feature interactions that stably affect the click over time, as shown by unvarying  $H^1 \rightarrow Y$  outside the box.
- $H^2$  denotes the feature interactions such as (stock, category) that affect the click differently over time due to the influence of time-varying environmental variable  $E$  (*e.g.*, the COVID-19 lockdown), as shown by time-varying  $E \rightarrow H^2 \rightarrow Y$  in the box.
- $H^3$  denotes the feature interactions that do not directly affect the click  $Y$  but are affected by  $E$  (differently over time), as shown by time-varying  $E \rightarrow H^3$  in the box.

Causally speaking,  $H^1$  has static causal relations to the click;  $H^2$  has dynamic causal relations to the click; and  $H^3$  is not causal feature interactions, but shows (spurious) correlations to the click, due to the confounding effect related to  $E$  (*e.g.*, the confounding effect brought by the backdoor path  $H^3 \leftarrow E \rightarrow H^2 \rightarrow Y$ ). As only static causal relations are invariant over time and exert stable effects on click generation, we treat  $H^1$  as stable feature interactions.

Nevertheless, the ERM paradigm exploits all correlations to fit the click data. The correlations between feature interactions and clicks could come from three sources: static causal relations ( $H^1 \rightarrow Y$ ), dynamic causal relations ( $H^2 \rightarrow Y$ ), and confounding effects (*e.g.*,  $H^3 \leftarrow E \rightarrow H^2 \rightarrow Y$ ). As a result, ERM inevitably learns unstable feature interactions, which have dynamic causal relations or spurious correlations to the click. It will make the model generalize poorly for future serving, since the effects of unstable feature interactions would vanish and even reverse [2, 9, 20, 44].

As such, we revise the objective of the CTR task as capturing stable feature interactions. Towards this goal, we split the historical data chronologically into several periods, enabling the learning of feature interactions that are stable across these periods. In the language of invariant learning [2], these periods form a set of environments, which are heterogeneous due to temporal influences, *e.g.*, the influence of the COVID-19 lockdown. Supposing the heterogeneity is large enough, only feature interactions with static causal relations to clicks are stably useful for CTR prediction across these environments. As such, a CTR model built in such a manner could learn stable feature interactions and generalize better to predict future clicks.

### 3.2 Disentangled Invariant Learning

In this subsection, we first discuss the potential and challenge of applying invariant learning to the CTR task, and then present the

*Disentangled Invariant Learning* (DIL) method from three aspects: optimization objective, learning algorithm, and model architecture.

**3.2.1 Invariant Learning Solution.** To achieve our goal, an ideal way is to conduct causal discovery, but it is hard in real-world scenarios due to the non-stationary, high-dimension, and partially unobservable challenges [10]. We propose to leverage invariant learning [2], which could capture the environment-invariant correlations given multiple heterogeneous environments, to learn stable feature interactions. There are two considerations: 1) periods could directly represent different environments, and 2) among the three types of correlations discussed in Section 3.1, the correlations brought by static causal relations are environment-invariant, while the correlations brought by dynamic causal relations and spurious correlations vary across environments<sup>2</sup>, *i.e.*, are environment-specific [20, 25]. It is thus possible to learn stable feature interactions by capturing environment-invariant correlations with invariant learning.

However, existing invariant learning methods are not applicable to the CTR task, because the sufficiency condition of their basic assumption (Assumption 1) is not satisfied.

- Clicks are affected by both feature interactions with static causal relations to clicks and feature interactions with dynamic causal relations to clicks, *i.e.*, clicks entangle environment-invariant and environment-specific correlations. That means clicks cannot be sufficiently predicted with only environment-invariant correlations, violating the sufficiency condition.

To overcome this challenge, the key lies in the disentanglement of environment-invariant and environment-specific correlations. With the disentanglement, the influence of environment-specific correlations could be eliminated from clicks, making the sufficiency condition satisfied. To this end, we propose a new disentangled invariant learning method for CTR models, which contains a new optimization objective, learning strategy, and model architecture.

**3.2.2 Disentangled Invariant Learning Objective.** We first divide the model parameters for feature interaction modeling into environment-invariant part  $\phi_s$  and environment-specific parts  $\{\phi_t\}_t$  (short for  $\{\phi_t\}_{t=1}^T$ , and  $\phi_t$  is used for  $t$ -th environment), to capture environment-invariant and environment-specific correlations, respectively. For example, we can equip the backbone model FM with environment-invariant feature embedding  $v_i^s$  and environment-specific feature embedding  $v_i^t$  for any  $i$ -th feature, and then  $\phi_s = \{v_i^s\}_i$  and  $\phi_t = \{v_i^t\}_i$  (*cf.* Section 3.2.4). With both  $\phi_s$  and  $\phi_t$ , we could sufficiently predict the click for a sample  $(x, y) \in \mathcal{D}_t$  as follows:

$$\hat{y} = f(x; \phi_s, \phi_t), \quad (5)$$

where  $f(x; \phi_s, \phi_t)$  is the modified CTR model. We present the detailed model architecture in Section 3.2.4 later.

To learn  $\phi_s$  and  $\{\phi_t\}_t$ , on one hand, we need an invariant learning objective, to make  $\phi_s$  capture the environment-invariant correlations. On the other hand, we need environment-specific objectives to force  $\{\phi_t\}_t$  to capture the environment-specific correlations. To this end, we propose the following overall optimization objective:

<sup>2</sup>Considering the correlations brought by  $H^2 \rightarrow Y$  and confounding effects (*e.g.*,  $H^3 \leftarrow E \rightarrow H^2 \rightarrow Y$ ), as  $E$  and the relations (*e.g.*,  $H^2 \rightarrow Y$ ) drift over time, these correlations vary over environments.

$$\min_{\phi_s} \sum_{t=1}^T w_t R(\mathcal{D}_t; \phi_s, \phi_t) + \lambda V_R \quad (6a)$$

$$s.t. \quad \{\phi_s, \phi_t\} \in \operatorname{argmin}_{\{\phi_s, \phi_t\}} R(\mathcal{D}_t; \phi_s, \phi_t) + \eta L_t, \forall t, \quad (6b)$$

where (6a) and (6b) correspond to the invariant learning objective and environment-specific objectives, respectively. We explain these learning objectives in detail below.

**- Invariant learning objective (Equation (6a)).** After  $\phi_t$  has captured the environment-specific correlations, fixing  $\phi_t$  is equal to removing its influences on clicks. Then capturing the environment-invariant correlations can sufficiently predict the clicks without the influence of environment-specific correlations, *i.e.*, the sufficiency condition can be satisfied. Thus we fix the environment-specific  $\phi_t$  to modify original invariant learning objective in Equation (3) for learning  $\phi_s$  as Equation (6a), in which  $R(\mathcal{D}_t; \phi_s, \phi_t)$  represents the loss on the  $\mathcal{D}_t$  collected from  $t$ -th environment and is computed with the *logloss* [24] as follows:

$$R(\mathcal{D}_t; \phi_s, \phi_t) = \sum_{(x, y) \in \mathcal{D}_t} -y \cdot \log(\sigma(f(x; \phi_s, \phi_t))) - (1 - y) \cdot \log(1 - \sigma(f(x; \phi_s, \phi_t))),$$

where  $\sigma(\cdot)$  denotes the sigmoid function. Besides, in Equation (6a),  $V_R$  refers to the variance of  $R(\mathcal{D}_t; \phi_s, \phi_t)$  over environments, computed like Equation (4); and  $w_t$  represents the weight for  $t$ -th environment, which is computed with the softmax function as follows:

$$w_t = \frac{e^{R(\mathcal{D}_t; \phi_s, \phi_t)}}{\sum_{t'=1}^T e^{R(\mathcal{D}_{t'}; \phi_s, \phi_{t'})}}, \quad (7)$$

where the temperature parameter of the softmax function is used but omitted here. As such, we assign a higher weight to the environment with higher loss, considering that paying more attention to difficult environments is helpful for improving the cross-environment performances [22, 33] to learn stable feature interaction.

**- Environment-specific objectives (Equation (6b)).** Combining  $\phi_s$  and  $\phi_t$  should fit  $\mathcal{D}_t$  well, since we need the clicks can be sufficiently predicted with  $\phi_s$  and  $\phi_t$ . Thus, for each environment  $t$ , we propose the optimization objective in Equation (6b), *i.e.*,

$$\min_{\phi_s, \phi_t} R(\mathcal{D}_t; \phi_s, \phi_t) + \eta L_t,$$

where  $\eta$  is a hyper-parameter to control the strength of  $L_t$ , and  $L_t$  is a regularizer to prevent  $\phi_t$  from capturing the environment-invariant correlations and is only used for learning  $\phi_t$ . Formally,

$$L_t = \sum_{t'=1, t' \neq t}^T [R(\mathcal{D}_{t'}; \phi_s, \phi_{t'}) - R(\mathcal{D}_{t'}; \phi_s, \phi_t)], \quad (8)$$

where  $R(\mathcal{D}_{t'}; \phi_s, \phi_t)$  represents the loss when predicting clicks in  $\mathcal{D}_{t'}$  forcibly using  $\phi_t$ . The key consideration is to make  $t$ -th environment-specific  $\phi_t$  not contribute to fitting data collected from other environments  $\{t' | 1 \leq t' \leq T, t' \in \mathbb{N}, t' \neq t\}$ .

The two optimization objectives are not isolated. The first objective is based on the second, since it requires  $\phi_t$  has captured environment-specific correlations. Thus we merge them to form the overall optimization objective in Equation (6). We term the overall optimization objective — disentangled invariant learning objective.

**3.2.3 Learning Strategy.** We now consider how to optimize the disentangled invariant learning objective (Equation (6)). As we need

**Algorithm 1:** DIL

---

**Input:** Historical data chronologically split into  $T$  environments, *i.e.*,  $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_t, \dots, \mathcal{D}_T\}$ .

- 1 Initialize  $\phi_s$  and  $\{\phi_t\}_t$ ;
- 2 **while** *Stop condition is not reached do*
- 3     Randomly sample a batch of data from  $\mathcal{D}$ ;
- 4     Normally update model parameters except  $\phi_s$  and  $\{\phi_t\}_t$ , and then keep them fixed;
- 5     Randomly sample an environment  $t \in \{1, \dots, T\}$ ;
- 6     // Update  $\phi_s$ , fixing all  $\phi_t$ ;
- 7     Compute  $\tilde{\phi}_s$  with Equation (9);
- 8     Update  $\phi_s$  according to Equation (10);
- 9     // Update  $\phi_t$ , fixing  $\phi_s$ ;
- 10    Update  $\phi_t$  according to Equation (11);
- 11 **end**

---

to remove the influences of environment-specific correlations on clicks by fixing  $\{\phi_t\}_t$  for the invariant learning objective, we propose a learning strategy, which updates the environment-invariant  $\phi_s$  and environment-specific  $\{\phi_t\}_t$  alternately as follows:

- **Update  $\phi_s$ .** Fixing  $\{\phi_t\}_t$ , we need to learn  $\phi_s$  such that it is optimal for both the invariant learning objective and environment-specific objective as shown in Equation (6). However, this is a hard bi-level optimization problem. Instead of directly solving it, we solve it in a meta-learning manner similar to [3]. Specially, we take the MAML [7] to update  $\phi_s$ , which has two main steps:

- Step 1. Meta training, which focuses on the environment-specific objective. We sample an environment  $t$  and compute the loss of environment-specific objective for it, then conduct an update to get a temporary  $\tilde{\phi}_s$  as follows:

$$\tilde{\phi}_s = \phi_s - \nabla_{\phi_s} R(\mathcal{D}_t; \phi_s, \phi_t). \quad (9)$$

- Step 2. Meta testing, which focuses on the invariant learning objective (Equation (6a)). We take the  $\tilde{\phi}_s$  to compute the loss of the invariant learning objective, and further update the  $\phi_s$  according to the loss. Formally,


$$\phi_s \leftarrow \phi_s - \nabla_{\phi_s} \left( \sum_{t'=1}^T w_{t'} R(\mathcal{D}_{t'}; \tilde{\phi}_s, \phi_{t'}) + \lambda V_R \right). \quad (10)$$

Note that  $V_R$  is also computed with  $\tilde{\phi}_s$ . In this way, it is easier to find a  $\phi_s$  that is generally good for the two optimization objectives, because of taking into account how an update for an environment-specific objective affects the invariant learning objective.

- **Update  $\phi_t$ .** After updating  $\phi_s$ , we directly optimize the environment-specific objective with  $\phi_s$  fixed, to update  $\phi_t$  (of the sampled environment) as follows:

$$\phi_t \leftarrow \phi_t - \nabla_{\phi_t} (R(\mathcal{D}_t; \phi_s, \phi_t) + \eta L_t). \quad (11)$$

These two updates are iterated alternately until convergence, and learning rates are used but omitted here. Algorithm 1 summarizes the detailed learning algorithm of DIL. In each iteration, we first randomly sample a batch of data and update all parameters except  $\phi_s$  and  $\phi_t$  with normal training loss (*e.g.*, logloss) on the data (lines 3-4). Then we randomly sample an environment  $t$  (line 5), and update  $\phi_s$  (lines 7-8). Last, we update  $\phi_t$  (lines 10). Note that, in each iteration, we only utilize the batch of data.



**Figure 3: Model Architectures.** (a) The embedding disentanglement model architecture for DIL. (b) The field-level disentanglement model architecture for LightDIL. "FI" denotes feature interactions. Red modules are environment-invariant, and the blue modules are environment-specific.

3.2.4 *Model Architecture.* We now consider designing the model architecture of  $f(x; \phi_s, \phi_t)$ , taking FM as the base CTR model.

- **Embedding disentanglement.** In light that the feature embedding is the core of feature interaction modeling (especially for FM-based models), we first directly conduct feature embedding disentanglement. Specifically, for any  $i$ -th feature, we let it have an environment-invariant embedding  $v_i^s$  and a set of environment-specific embeddings  $\{v_i^t\}_t$ , as shown in Figure 3(a). Then for a sample  $(x, y) \in \mathcal{D}_t$  belonging to  $t$ -th environment, we modify FM (Equation (1)) to generate the prediction as follows:

$$\hat{y} = f(x; \phi_s, \phi_t) = \sum_{i=1}^N \sum_{j>i}^N \langle v_i^s + v_i^t, v_j^s + v_j^t \rangle \cdot x_i x_j, \quad (12)$$

where  $\phi_s = \{v_i^s\}_i$  and  $\phi_t = \{v_i^t\}_i$ . This is the default disentanglement for the proposed Disentangled Invariant Learning (DIL)<sup>3</sup>.

- **Field-level disentanglement.** The embedding disentanglement would result in  $T$  times more model parameters since each environment corresponds to a set of environment-specific embeddings, which dominates the model parameters in FM. What is worse, it is hard to train a model with massive parameters well. This drives us to design a light version of DIL. Note that existing methods [24, 28] assign weights for feature interactions to model their importance, *i.e.*, the correlation strength to the click. We could similarly assign environment-invariant and environment-specific weights for feature interactions, to capture the environment-invariant and environment-specific correlations, respectively. However, partial feature interactions could appear only in the future but not appear in history [32], resulting in their weights being unlearnable. To overcome the issue, we assign weights at the field level, as shown in Figure 3(b). Then we modify the FM in Equation (2) as follows:

$$\hat{y} = f(x; \phi_s, \phi_t) = \sum_{i=1}^M \sum_{j>i}^M (\alpha_{ij}^s + \alpha_{ij}^t) \langle u_i, u_j \rangle, \quad (13)$$

where  $\alpha_{ij}^s \in \mathbb{R}$  ( $\alpha_{ij}^t \in \mathbb{R}$ ) denotes the environment-invariant ( $t$ -th environment-specific weight) for the feature interaction between field  $i$  and  $j$ ,  $\phi_s = \{\alpha_{ij}^s\}_{ij}$ , and  $\phi_t = \{\alpha_{ij}^t\}_{ij}$ . Note that

<sup>3</sup>There is another consideration for designing embedding disentanglement – it is widely applicable since CTR models usually project features into embeddings.

feature embeddings do not belong to  $\phi_s$  and  $\phi_t$  here, and are normally updated (see Algorithm 1). We name the DIL performing the disentanglement at the higher level of the feature field LightDIL.

- **Inference.** During inference, DIL and LightDIL only take  $\phi_s$  to generate predictions, to pursue good generalization performance. For example, LightDIL generates the prediction as:  $\hat{y} = f(x; \phi_s, \emptyset) = \sum_{i=1}^M \sum_{j>i}^M \alpha_{ij}^s \langle \mathbf{u}_i, \mathbf{u}_j \rangle$ . Another way is combining the  $\phi_s$  learned in training and the  $\phi_t$  estimated for a future period to generate predictions. However, it is hard to forecast the  $\phi_t$  of a future period. We leave it as future work.

## 4 EXPERIMENTS

We conduct experiments to answer three research questions:

**RQ1:** How does DIL perform on real-world data? **RQ2:** How do the design choices of DIL affect its performance? **RQ3:** Can DIL learn stable feature interactions when there are spurious correlations? And can DIL learn stable feature interactions when there are dynamic causal relations (*i.e.*, the sufficiency condition of invariant learning is not satisfied)?

### 4.1 Experimental Settings

• **Baselines.** We compare the proposed methods with the following recommender methods: 1) the basic FM [32]; 2) FwFMs [28], which models the importance of feature interactions via field weights; 3) AFM [42], which utilizes an attention network to learn the importance of feature interactions; 4) AutoFIS [24] and PROFIT [8], which are advanced neural architecture search (NAS) based methods for selecting more effective feature interactions; 5) CFM [20], which is a method for learning robust feature interactions (defined as causal feature interactions); In addition, we further compare DIL with two Out-of-Distribution (OOD) generalization methods, including one representative invariant learning method – V-REx [18] and a distributionally robust optimization method – Group-DRO [33]. For a fair comparison, we implement V-REx and Group-DRO based on FM. Considering that we aim at learning stably useful feature interactions to generalize well for the future, we select baselines 1) modeling the usefulness of feature interactions (FwFMs, AFM, AutoFIS, PROFIT, and CFM), or 2) pursuing the capability of OOD generalization (V-REx, Group-DRO, and CFM).

• **Evaluation metrics and hyper-parameters.** The widely used metrics for CTR prediction are Area under the ROC Curve (AUC) and logloss (*i.e.*, binary cross entropy). In this work, to evaluate whether a model can generalize well to multiple future periods (testing environments), we report the average values of AUC and logloss over different testing environments.

For a fair comparison, we learn all models based on the binary cross entropy loss and tune them according to the AUC metric computed on the validation environments. Meanwhile, for all models, we set the feature embedding size as 48 and take the Adam [17] as the optimizer with the default batch size of 8,192. We leverage the grid search to find the best hyper-parameters for each model. In particular, we search the learning rate in the range of  $\{1e-2, 1e-3, 1e-4\}$ , and the  $L_2$  regularization coefficient of the embedding layers in the range of  $\{1e-1, 1e-2, \dots, 1e-7\}$  for all models. For the special hyper-parameters of baselines, we search most of them in the ranges provided by their papers. For the special hyper-parameters  $\lambda$  and  $\eta$  of our method (in Equation (6)), we search both

**Table 1: Statistics of the evaluation datasets.**

Dataset	#Instances	#Fields	#Features	#Train	#Valid	#Test
ML-10M	$1.0 \times 10^7$	7	145896	$1.7 \times 10^6$	$1.9 \times 10^6$	$1.5 \times 10^6$
Douban	$4.2 \times 10^6$	19	307919	$7.9 \times 10^5$	$6.8 \times 10^5$	$1.4 \times 10^6$
Avazu	$4.0 \times 10^7$	23	645185	$2.0 \times 10^7$	$7.1 \times 10^6$	$1.3 \times 10^7$

of them in the range of  $\{1e-1, \dots, 1e-4\}$ . We release the code at <https://github.com/zyang1580/DIL>.

### 4.2 Experiments on Real-world Data

**4.2.1 Real-world Datasets.** We conduct extensive experiments on the two representative datasets for CTR prediction: MovieLens-10M and Douban, the statistics of which are summarized in Table 1.

- **MovieLens-10M (ML-10M)** [13] is released by GroupLens Research, containing ratings collected from 1995 to 2009 with rich features belonging to various fields, such as age, gender, and movie category. We only preserve the data collected from July 2002 to December 2008, considering the data sparsity of other years. We then chronologically split the data into 13 periods by treating six months as a period. We further divide these periods into the training, validation, and testing environments according to the ratio of 5:4:4 with the order kept.

- **Douban**<sup>4</sup> is a popular dataset for CTR prediction, which contains the rating data collected from 2012 to 2019 on Douban Website. It contains user and item features belonging to the field of language, actors, *etc.* Similarly, we preserve the data ranging from January 2012 to June 2019 and then split it into 15 periods with the first 5 for training, the middle 5 for validation, and the last 5 for testing.

In both ML-10M and Douban, the ratings range from 1 to 5. Following [38], we treat the user-item interactions with ratings  $\geq 3$  as positive samples with the click label  $y = 1$  (otherwise,  $y = 0$ ). Meanwhile, to ensure the dataset quality and avoid inadequate learning for extremely sparse features, we set the features appearing less than 10 times as a dummy feature ('other') [8, 24].

**4.2.2 Overall Performance Comparison (RQ1).** Table 2 summarizes the overall performance comparison between the baselines, DIL, and LightDIL on Douban and ML-10M datasets. From the table, we have the following observations.

- LightDIL achieves the best performance on both ML-10M and Douban. DIL is slightly worse than LightDIL but still outperforms all baselines on most of the metrics. Compared with the FM, FwFMs, AFM, AutoFIS, and PROFIT, which are learned with ERM, the better results of DIL indicate that learning stable feature interactions could bring better generalization for future serving than pursuing ERM on historical data. The results verify 1) the rationality of our reformulation for the CTR task and 2) the effectiveness of our methods in learning stable feature interactions. Compared to V-REx and Group-DRO, better results of DIL verify that DIL can overcome the challenge – the clicks cannot be sufficiently predicted with environment-invariant correlations – to capture environment-invariant correlations, learning stable feature interactions in CTR prediction.
- FwFMs, AFM, AutoFIS, and PROFIT, which learn more effective and discard less useful feature interactions, outperform FM on Douban, but fail to surpass FM and even cause worse performance on ML-10M<sup>5</sup>. These models learn feature interactions with strong

<sup>4</sup><https://www.csuldw.com/2019/09/08/2019-09-08-moviedata-10m/>.


<sup>5</sup>AFM and PROFIT beat FM on MovieLens in their papers, which are different from our results. This is because their training and testing data are identically distributed.



**Table 2: Performance comparison between the baselines, DIL, and LightDIL. "Rel. Impr." denotes the relative improvement over FM *w.r.t.* AUC. The value on the left (right) of "±" represents the average (the standard deviation) *w.r.t.* the corresponding metric.**

Model	Douban			ML-10M		
	AUC ↑	logloss ↓	Rel. Impr.	AUC ↑	logloss ↓	Rel. Impr.
FM	0.7959 ± 0.0076	0.4459 ± 0.0269	-	0.7158 ± 0.0081	0.4439 ± 0.0202	-
FwFMs	0.8071 ± 0.0062	0.4262 ± 0.0217	1.41%	0.7146 ± 0.0081	0.4465 ± 0.0202	-0.17%
AFM	0.8055 ± 0.0056	0.4271 ± 0.0207	1.21%	0.7095 ± 0.0081	0.4579 ± 0.0167	-0.88%
AutoFIS	0.8064 ± 0.0059	0.4247 ± 0.0201	1.32%	0.7135 ± 0.0074	0.4488 ± 0.0208	-0.32%
PROFIT	0.8077 ± 0.0069	0.4248 ± 0.0216	1.48%	0.7158 ± 0.0079	0.4498 ± 0.0216	0
CFM	0.8023 ± 0.0060	0.4289 ± 0.0208	0.80%	0.7091 ± 0.0088	0.4473 ± 0.0151	-0.93%
Group-DRO	0.8006 ± 0.0060	0.4290 ± 0.0205	0.59%	0.7133 ± 0.0081	0.4443 ± 0.0186	-0.35%
V-REx	0.7996 ± 0.0065	0.4313 ± 0.0223	0.46%	0.7112 ± 0.0081	0.4487 ± 0.0199	-0.64%
DIL	0.8081 ± 0.0058	0.4230 ± 0.0193	1.53%	0.7162 ± 0.0082	0.4447 ± 0.0195	0.05%
LightDIL	<b>0.8089 ± 0.0057</b>	<b>0.4219 ± 0.0193</b>	1.63%	<b>0.7171 ± 0.0076</b>	<b>0.4438 ± 0.0198</b>	0.18%

We have conducted a *t*-test, verifying the improvements of our best method to the best baseline are statistically significant on AUC (p-value < 0.05).

**Figure 4: Comparison between FM, DIL, and LightDIL in terms of model size (left) and training time (right).**

correlations to clicks in historical data. They thus are possibly dominated by spurious correlations while ignoring the causal relations, leading to poor generalization performance. ML-10M has fewer fields than Douban as shown in Table 1, and thus the detrimental effect of mistakenly discarding causally useful feature interactions is more severe on ML-10M.

- V-REx and Group-DRO outperform FM on Douban, but not on ML-10M, verifying that directly applying OOD generation methods for recommendation is not effective. For the invariant learning method V-REx, its basic assumption (Assumption 1) cannot be satisfied in the CTR task. Group-DRO just focuses on the worst environment, sacrificing the overall performance [29].
- CFM, which aims to learn personalized causal feature interactions, shows inferior performance. We attribute its poor performance to inaccurate causal effects estimation and capturing dynamic causal relations. On one hand, for de-confounding, CFM needs to learn balancing weights for all training samples, which is non-trivial due to the large size of our datasets (at least 50 times larger than the dataset used in the CFM paper). And it is hard to estimate the personalized causal effects well due to the sparsity of user data. On the other hand, CFM captures dynamic causal relations, which could lead to poor generalization performances.

**4.2.3 In-depth Analyses (RQ2).** In this part, we first compare FM, DIL, and LightDIL on two datasets *w.r.t.* model size and training cost. We then conduct ablation studies on LightDIL to validate the effectiveness of different components of our method. Last, we study how the granularity of environment splitting affects the performance of our proposal based on LightDIL.

• **Modeling efficiency.** We conduct a comparison between FM, DIL, and LightDIL regarding the model size and training time. As shown in Figure 4, LightDIL has almost the same model size as


FM, which is far smaller than that of DIL. Meanwhile, LightDIL could greatly reduce the time cost of training, compared to DIL. The results verify the validity of the field-level disentanglement in improving modeling efficiency. Moreover, LightDIL improves the CTR performance over DIL, as illustrated in Table 2. It reflects that the field-level disentanglement is sufficient to distinguish stable and unstable feature interactions with much fewer parameters. Regarding the performance difference between LightDIL and DIL, we attribute it to the discrepancy between them *w.r.t.* the disentanglement. For example, embedding disentanglement brings more model parameters for DIL while the corresponding feature interactions are more sparse (compared to the interaction between feature fields), making DIL hard to learn well. Regarding training, LightDIL is about five times slower than FM. This is because we adopt the paradigm of MAML [7] to update environment-invariant model parameters, which costs highly when computing the second-order gradients. In the future, it is possible to speed up LightDIL with the first-order approximation for MAML [7].

• **Ablation studies.** We next conduct ablation experiments to validate the effectiveness of the environment weights  $w_t$  in Equation (6), the variance-based regularizer  $V_R$  in Equation (6), the meta-learning module, the disentanglement, and the environment-specific regularizer  $L_t$  in Equation (8). Table 3 enumerates the performance of the variants of LightDIL by disabling the above components one by one. From the table, we obtain the following findings.

- LightDIL w/o environment weights replaces the weights  $w_t$  in Equation (6) with the uniform  $\frac{1}{T}$ . In Table 3, the performance gap between LightDIL and the variant w/o environment weights shows that using  $w_t$  instead of equally treating each environment is helpful. The underlying reason is that using the weights will lead the model to focus more on the difficult environments during the optimization, validating the arguments in Section 3.2.2.
- The performance of LightDIL drops when we discard the variance-based regularizer, reflecting its importance to balance the model's predictions over multiple environments. The variance-based regularizer can regulate the model to be robust across environments, helping to learn the environment-invariant correlations.
- The variant w/o meta-learning is implemented by removing meta training in Equation (9) and replacing  $\hat{\phi}_s$  with  $\phi_s$  in Equation (10) to update  $\phi_s$ . The inferior results of this variant indicate that meta-learning does help solve the challenging bi-level optimization problem in Equation (6) and enhance the CTR performance.

**Table 3: Results of the ablation studies over LightDIL on Douban and ML-10M.**

Model	Douban		ML-10M	
	AUC $\uparrow$	logloss $\downarrow$	AUC $\uparrow$	logloss $\downarrow$
w/o environment weights	0.8069 $\pm$ 0.0057	0.4233 $\pm$ 0.0200	0.7156 $\pm$ 0.0078	0.4451 $\pm$ 0.0198
w/o variance-based regularizer	0.8075 $\pm$ 0.0060	0.4243 $\pm$ 0.0212	0.7140 $\pm$ 0.0074	0.4479 $\pm$ 0.0210
w/o meta-learning	0.8057 $\pm$ 0.0065	0.4250 $\pm$ 0.0213	0.7122 $\pm$ 0.0084	0.4454 $\pm$ 0.0148
w/o disentanglement	0.8068 $\pm$ 0.0058	0.4251 $\pm$ 0.0210	0.7107 $\pm$ 0.0085	0.4542 $\pm$ 0.0128
w/o environment-specific regularizer	0.8062 $\pm$ 0.0056	0.4255 $\pm$ 0.0203	0.7150 $\pm$ 0.0079	0.4439 $\pm$ 0.0190
LightDIL (original)	0.8089 $\pm$ 0.0057	0.4219 $\pm$ 0.0193	0.7171 $\pm$ 0.0076	0.4438 $\pm$ 0.0198


**Figure 5: Performance of LightDIL on Douban (left) and ML-10M (right) when splitting training data into different numbers of environments.**

- Comparing LightDIL with the variant w/o disentanglement, we can observe that the disentanglement module effectively improves the CTR accuracy across environments. The better result of LightDIL shows that the disentanglement is vitally important for capturing the environment-invariant correlations when the sufficiency condition (in Assumption 1) is not satisfied.
- We can find that the performance decreases when the environment-specific regularizer  $L_t$  is disabled. The result verifies that  $L_t$  could force  $\phi_t$  to capture environment-specific correlations, and it is important for the disentanglement of  $\phi_s$  and  $\phi_t$ .
- **Granularity of separating environments.** We next study how the granularity of environment splitting affects the effectiveness of our proposal. We re-split the training data into  $T$  environments and evaluate the performance of the LightDIL trained on the new split environments. We draw the performance curve of LightDIL when  $T$  varies in  $\{3, 5, 7, 9, 11, 15\}$ . Figure 5 summarizes the results. We find that the performance of LightDIL increases first and then decreases as  $T$  increases, *i.e.*, the granularity becomes small. We attribute it to that a more fine-grained splitting could increase the heterogeneity of training environments, benefiting the learning of the environment-invariant correlations. However, too fine-grained splitting would decrease the quality of training environments, *e.g.*, the data becomes very sparse, preventing the exhibition of environment-invariant correlations. Besides, we find LightDIL could beat the best baseline when  $T$  varies in a wide range, verifying the superiority of our method.

### 4.3 Experiments on Semi-synthetic Data (RQ3)

We further design experiments on semi-synthetic datasets to verify whether DIL can successfully capture stable feature interactions by discarding the two types of environment-specific correlations: spurious correlations and correlations brought by dynamic causal relations, respectively. Note that if dynamic causal relations exist, the sufficiency condition (Assumption 1) is not satisfied.

• **Semi-synthetic datasets.** It is hard to distinguish spurious correlations and dynamic causal relations in real-world datasets. As such, we construct two semi-synthetic datasets based on **Avazu**,

**Figure 6: Performance of DIL, LightDIL, and the baselines on two semi-synthetic datasets: SP-Avazu and DC-Avazu.**

which is an advertising click dataset provided by Avazu corporation in the Kaggle CTR Prediction Contest<sup>6</sup>. The statistics of Avazu are provided in Table 1. Avazu contains the click data between users and items spanning 10 days. We aim to separately inject the spurious correlations and dynamic causal relations into two constructed datasets. And thus we hope the feature interactions in the original dataset are relatively stable. With this consideration, we select Avazu with the click data in a short period, where the correlations between the original feature interactions and the clicks are unlikely to shift. Thereafter, we treat each day as an environment and construct two datasets: Avazu with spurious correlations (SP-Avazu) and Avazu with dynamic causal relations (DC-Avazu)<sup>7</sup>. For both SP-Avazu and DC-Avazu, we allocate the first 5 environments for training, the middle 2 for validation, and the last 3 for testing.

• **Performance comparison.** We conduct experiments on both SP-Avazu and DC-Avazu to verify that DIL could learn stable feature interactions with the existence of strong spurious correlations and dynamic causal relations. We compare our methods DIL and LightDIL with FM, FwFMs, AFM, AutoFIS, Group-DRO, and V-REx. Here we ignore PROFIT and only compare with one representative NAS-based model AutoFIS because of the extremely long searching time of NAS-based models. Besides, we do not compare with CFM due to the lack of user ID information on Avazu, which is necessary for CFM. Figure 6 illustrates the performance of these methods. We can draw several conclusions from the figure.

- On SP-Avazu, DIL and LightDIL substantially outperform the baselines. Moreover, they can generate meaningful predictions (AUC > 0.5) while most baselines cannot achieve that. These results verify that the proposed disentangled invariant learning method can learn stable feature interactions and discard the feature interactions with spurious correlations.
- On SP-Avazu, FM, FwFMs, AFM, and AutoFIS show poor performance with AUC close to or less than 0.5<sup>8</sup>. The results reflect that models built with the ERM paradigm are significantly affected by the feature interactions with strong spurious correlations.

<sup>6</sup> <http://www.kaggle.com/c/avazu-ctr-prediction>.

<sup>7</sup> The details could be found in the Arxiv version of the paper.

<sup>8</sup> AUC < 0.5 due to inverse spurious correlations of training and testing environments.



- On DC-Avazu, DIL and LightDIL also outperform all the baselines. Especially, DIL and LightDIL achieve better performances than V-REx on DC-Avazu. This shows the effectiveness of the proposed method in learning stable feature interactions in the case where the sufficiency condition in Assumption 1 is not satisfied. On SP-Avazu, we unexpectedly find that V-REx performs very poorly. The possible reason is that V-REx needs large enough distribution shifts between environments, which is not satisfied. The spurious correlations in training environments are mostly positive and thus the shifts might not be large enough. In contrast, the superior performance of DIL indicates that DIL can still learn stable feature interactions from the training environments with small shifts.
- Regarding Group-DRO, it achieves good performance on SP-Avazu but performs poorly on DC-Avazu. Group-DRO focuses on the worst-case environment, which helps to avoid being overly reliant on spurious correlations. Nevertheless, ignoring some environments makes Group-DRO possibly capture the environment-specific correlations, bringing bad performance on DC-Avazu.

To sum up, DIL and LightDIL can capture stable feature interactions and discard unstable feature interactions even when the sufficiency condition of invariant learning is not satisfied, enabling them to generalize well to unknown future periods.

## 5 RELATED WORK

**CTR prediction.** CTR prediction has been widely studied for many years [43]. Since raw features rarely lead to satisfying results [27, 43], feature interaction modeling, which indicates the combination relationships of multiple features, becomes the focus of CTR prediction. FM is a pioneer and classical method, which models feature interactions in a factorization manner [32]. Later, many methods are proposed to achieve more effective and complicated feature interaction modeling based on various neural networks, such as MLPs [4, 11], product-based neural networks [14, 30], attention-based networks [34, 42], convolutional neural networks [23, 26], and graph neural networks [12, 19]. Recently, NAS-based methods are proposed to automatically search the optimal network architecture for feature interaction modeling [46], and automatically select/generate more effective feature interactions [8, 24, 37]. These methods reduce human efforts and achieve better feature interaction modeling. Besides, some methods [39, 45] additionally consider sequential user behavior modeling.

Although various CTR models are proposed, existing works are built in an ERM paradigm, blindly learning feature interactions useful for fitting historical clicks. On the contrary, we reformulate the CTR task to learn stable feature interactions and propose an invariant learning solution. We notice that CFM [20] is also not built with ERM. Differently, CFM aims to capture all causal feature interactions via distribution balancing, while we learn stable feature interactions (with static causal relations to click) instead of all causal feature interactions, using the invariant learning method.

**Invariant learning.** To alleviate spurious correlations, previous studies have focused on invariant learning [2, 25]. Generally speaking, invariant learning assumes that the training data is collected from distinct environments and aims to pursue robust predictions across multiple environments. In particular, Invariant Risk Minimization (IRM) [2] discovers stable features from multiple environments. Following IRM, some studies pay attention to relaxing

the linear assumption in IRM [1], automatically splitting environments [5, 25], and revising the regularizers [18]. These methods assume the target could be sufficiently predicted with environment-invariant correlations, which is not satisfied in the CTR task. We extend invariant learning to overcome this challenge. Except for invariant learning solutions, Distributionally Robust Optimization improves the OOD generalization by minimizing the loss in the worst environment [21, 33]. Stable learning instead achieves robust predictions via confounder balancing [35].

Existing works have also considered invariant learning for recommendation debiasing [15, 40], and alleviating spurious correlations in multimedia recommendation [6]. SGL [15] focuses on dealing with selection bias. InvPref [40] infers the labels of heterogeneous environments and captures invariant preference across environments for debiasing in collaborative filtering. InvRL [6] takes a similar method to InvPref [40] but aims at alleviating spurious correlations in multimedia content. Different from our proposal, these methods are not designed for the CTR task, do not learn stable feature interactions, and do not overcome the challenge that the sufficiency condition is not satisfied. Besides, existing works have also considered achieving OOD generalization for recommendation by learning causal relations (static and dynamic) [16, 20, 38]. COR [38] and CausalPref [16] both focus on learning the causal model of data generation. CFM [20] is a CTR model and achieves OOD recommendations by distribution balancing as discussed before. Different from them, we aim at learning stable feature interactions with **static** causal relations to clicks.

## 6 CONCLUSION

In this work, we reformulate the CTR prediction for recommendation by learning stable feature interaction from split environments, aiming to generalize well for future serving. Towards this goal, we propose a novel disentangled invariant learning framework, which extends invariant learning to successfully capture invariant correlations when the prediction target entangles both environment-invariant and environment-specific correlations, *i.e.*, the sufficiency prediction condition is not satisfied. We conduct extensive experiments on both real-world and semi-synthetic datasets, providing insightful analysis for the effectiveness of our proposal.

This work showcased the limitation of the ERM learning paradigm for CTR prediction, despite its dominant role in recommendation research and industry. We replace the ERM learning paradigm with the proposed disentangled invariant learning to learn stable information, achieving better generalization for future periods. We believe the proposed invariant learning paradigm could adapt to other real-world recommendation tasks and recommender models. In the future, we will apply our DIL to other state-of-the-art recommender models and other recommendation tasks (*e.g.*, collaborative filtering, sequential recommendation, and conversational recommendation [31]). Besides, we will also explore better disentanglement mechanisms for the proposed DIL.

## ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China (2021YFF0901603), the National Natural Science Foundation of China (62272437, 62121002), and the CCCD Key Lab of Ministry of Culture and Tourism.

## REFERENCES

- [1] Kartik Ahuja, Karthikeyan Shanmugam, Kush Varshney, and Amit Dhurandhar. 2020. Invariant Risk Minimization Games. In *International Conference on Machine Learning*. PMLR, 145–155.
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893* (2019).
- [3] Jun-Hyun Bae, Inchul Choi, and Minho Lee. 2021. Meta-Learned Invariant Risk Minimization. *arXiv preprint arXiv:2103.12947* (2021).
- [4] Weijie Bian, Kailun Wu, Lejian Ren, Qi Pi, Yujing Zhang, Can Xiao, Xiang-Rong Sheng, Yong-Nan Zhu, Zhangming Chan, Na Mou, et al. 2022. CAN: Feature Co-Action Network for Click-Through Rate Prediction. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 57–65.
- [5] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. 2021. Environment Inference for Invariant Learning. In *International Conference on Machine Learning*. 2189–2200.
- [6] Xiaoyu Du, Zike Wu, Fuli Feng, Xiangnan He, and Jinhui Tang. 2022. Invariant Representation Learning for Multimedia Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*. 619–628.
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*. 1126–1135.
- [8] Chen Gao, Yinfeng Li, Quanming Yao, Depeng Jin, and Yong Li. 2021. Progressive Feature Interaction Search for Deep Sparse Network. *Advances in Neural Information Processing Systems* 34 (2021).
- [9] Chen Gao, Yu Zheng, Wenjie Wang, Fuli Feng, Xiangnan He, and Yong Li. 2022. Causal Inference in Recommender Systems: A Survey and Future Directions. *arXiv preprint arXiv:2208.12397* (2022).
- [10] Clark Glymour, Kun Zhang, and Peter Spirtes. 2019. Review of causal discovery methods based on graphical models. *Frontiers in genetics* 10 (2019), 524.
- [11] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. 1725–1731.
- [12] Wei Guo, Rong Su, Renhao Tan, Huifeng Guo, Yingxue Zhang, Zhirong Liu, Ruiming Tang, and Xiuqiang He. 2021. Dual graph enhanced embedding neural network for ctr prediction. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 496–504.
- [13] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm Transactions on Interactive Intelligent Systems* 5, 4 (2015), 1–19.
- [14] Xiangnan He and Tat-Seng Chua. 2017. Neural factorization machines for sparse predictive analytics. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 355–364.
- [15] Yue He, Peng Cui, Jianxin Ma, Hao Zou, Xiaowei Wang, Hongxia Yang, and Philip S Yu. 2020. Learning Stable Graphs from Multiple Environments with Selection Bias. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2194–2202.
- [16] Yue He, Zimu Wang, Peng Cui, Hao Zou, Yafeng Zhang, Qiang Cui, and Yong Jiang. 2022. CausPref: Causal Preference Learning for Out-of-Distribution Recommendation. In *Proceedings of the ACM Web Conference 2022*. 410–421.
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Remi Le Priol, and Aaron Courville. 2021. Out-of-distribution Generalization via Risk Extrapolation (rex). In *International Conference on Machine Learning*. PMLR, 5815–5826.
- [19] Feng Li, Zhenrui Chen, Pengjie Wang, Yi Ren, Di Zhang, and Xiaoyu Zhu. 2019. Graph intention network for click-through rate prediction in sponsored search. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 961–964.
- [20] Yunqi Li, Hanxiong Chen, Juntao Tan, and Yongfeng Zhang. 2022. Causal factorization machine for robust recommendation. In *Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries*. 1–9.
- [21] Fengming Lin, Xiaolei Fang, and Zheming Gao. 2022. Distributionally robust optimization: A review on theory and applications. *Numerical Algebra, Control and Optimization* 12, 1 (2022), 159–212.
- [22] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision*. 2980–2988.
- [23] Bin Liu, Ruiming Tang, Yingzhi Chen, Jinkai Yu, Huifeng Guo, and Yuzhou Zhang. 2019. Feature Generation by Convolutional Neural Network for Click-Through Rate Prediction. In *The World Wide Web Conference*. 1119–1129.
- [24] Bin Liu, Chenxu Zhu, Guilin Li, Weinan Zhang, Jincai Lai, Ruiming Tang, Xiuqiang He, Zhenguo Li, and Yong Yu. 2020. Autofis: Automatic feature interaction selection in factorization models for click-through rate prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2636–2645.
- [25] Jiashuo Liu, Zheyuan Hu, Peng Cui, Bo Li, and Zheyuan Shen. 2021. Heterogeneous Risk Minimization. In *International Conference on Machine Learning*. 6804–6814.
- [26] Qiang Liu, Feng Yu, Shu Wu, and Liang Wang. 2015. A convolutional click prediction model. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. 1743–1746.
- [27] Erxue Min, Yu Rong, Tingyang Xu, Yatao Bian, Da Luo, Kangyi Lin, Junzhou Huang, Sophia Ananiadou, and Peilin Zhao. 2022. Neighbour Interaction based Click-Through Rate Prediction via Graph-masked Transformer. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 353–362.
- [28] Junwei Pan, Jian Xu, Alfonso Lobos Ruiz, Wenliang Zhao, Shengjun Pan, Yu Sun, and Quan Lu. 2018. Field-Weighted Factorization Machines for Click-Through Rate Prediction in Display Advertising. In *Proceedings of the 2018 World Wide Web Conference*. 1349–1357.
- [29] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. 2022. Focus on the Common Good: Group Distributional Robustness Follows. In *International Conference on Learning Representations*.
- [30] Yanru Qu, Bohui Fang, Weinan Zhang, Ruiming Tang, Minzhe Niu, Huifeng Guo, Yong Yu, and Xiuqiang He. 2018. Product-based neural networks for user response prediction over multi-field categorical data. *ACM Transactions on Information Systems (TOIS)* 37, 1 (2018), 1–35.
- [31] Zhaochun Ren, Zhi Tian, Dongdong Li, Pengjie Ren, Liu Yang, Xin Xin, Huasheng Liang, Maarten de Rijke, and Zhumin Chen. 2022. Variational Reasoning about User Preferences for Conversational Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 165–175.
- [32] Steffen Rendle. 2010. Factorization machines. In *2010 IEEE International Conference on Data Mining*. IEEE, 995–1000.
- [33] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally Robust Neural Networks. In *International Conference on Learning Representations*.
- [34] Arindam Sarkar, Dipankar Das, Vivek Sembium, and Prakash Mandayam Comar. 2022. Dual Attentional Higher Order Factorization Machines. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 378–388.
- [35] Zheyuan Shen, Peng Cui, Kun Kuang, Bo Li, and Peixuan Chen. 2018. Causally Regularized Learning with Agnostic Data Selection Bias. In *Proceedings of the 26th ACM International Conference on Multimedia*. 411–419.
- [36] Zheyuan Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. 2021. Towards Out-of-distribution Generalization: A Survey. *arXiv preprint arXiv:2108.13624* (2021).
- [37] Yixin Su, Yunxiang Zhao, Sarah Erfani, Junhao Gan, and Rui Zhang. 2022. Detecting Arbitrary Order Beneficial Feature Interactions for Recommender Systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1676–1686.
- [38] Wenjie Wang, Xinyu Lin, Fuli Feng, Xiangnan He, Min Lin, and Tat-Seng Chua. 2022. Causal Representation Learning for Out-of-Distribution Recommendation. In *Proceedings of the ACM Web Conference 2022*. 3562–3571.
- [39] Yifan Wang, Yifang Qin, Fang Sun, Bo Zhang, Xuyang Hou, Ke Hu, Jia Cheng, Jun Lei, and Ming Zhang. 2022. DisenCTR: Dynamic Graph-based Disentangled Representation for Click-Through Rate Prediction. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2314–2318.
- [40] Zimu Wang, Yue He, Jiashuo Liu, Wenchao Zou, Philip S Yu, and Peng Cui. 2022. Invariant Preference Learning for General Debiasing in Recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1969–1978.
- [41] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. 2022. Handling Distribution Shifts on Graphs: An Invariance Perspective. In *International Conference on Learning Representations*.
- [42] Jun Xiao, Hao Ye, Xiangnan He, Hanwang Zhang, Fei Wu, and Tat-Seng Chua. 2017. Attentional Factorization Machines: Learning the Weight of Feature Interactions via Attention Networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. 3119–3125.
- [43] Weinan Zhang, Jiarui Qin, Wei Guo, Ruiming Tang, and Xiuqiang He. 2021. Deep Learning for Click-Through Rate Estimation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. 4695–4703.
- [44] 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.
- [45] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1059–1068.
- [46] Guanghui Zhu, Feng Cheng, Defu Lian, Chunfeng Yuan, and Yihua Huang. 2022. NAS-CTR: Efficient Neural Architecture Search for Click-Through Rate Prediction. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 332–342.