

IdeFN: Identifying Unclicked Space False Negatives via Relaxed Partial Optimal Transport for Conversion Rate Prediction

Weiye Zhong¹, Weiming Liu², Lianyong Qi^{1,3*}, Xiaoran Zhao³, Xiaolong Xu⁴, Haolong Xiang⁴, Yang Cao⁵, Shichao Pei⁶, Qiang Ni⁷

¹Qufu Normal University, China

²ByteDance Inc., Singapore

³China University of Petroleum (East China), China

⁴Nanjing University of Information Science and Technology, China

⁵Great Bay University, China

⁶University of Massachusetts Boston, USA

⁷Lancaster University, United Kingdom

weiyizhong@qfnu.edu.cn, {lwming95, lianyongqi}@gmail.com, xiaoranzhao@s.upc.edu.cn, {xllxu, hlxiang}@nuist.edu.cn, charles.cao@ieee.org, shichao.pei@umb.edu, q.ni@lancaster.ac.uk

Abstract

Accurate conversion rate (CVR) prediction is critical for recommender systems to capture user conversion intent and increase platform revenues. Traditional CVR models commonly suffer from sample selection bias (SSB) and data sparsity (DS), which has led to the adoption of click-through & conversion rate (CTCVR) multi-task learning frameworks to alleviate these issues. However, existing methods implicitly mislabel some unclicked samples with genuine conversion potential as negatives, thereby exacerbating the false negative sample (FNS) problem. To address this, we propose IdeFN, a multi-task CVR framework that identifies false negatives in the unclicked space to enable CVR prediction across the entire exposure space and leverages click-through rates (CTR) as an auxiliary task for shared-parameter learning. Specifically, IdeFN consists of two main components, i.e., relaxed partial optimal transport (RPOT) module and sample relabeling mechanism (SRM). The former estimates the soft matching strengths between unclicked samples and positive samples under a relaxed partial optimal transport formulation, establishing corresponding relationships between these samples. The latter adaptively re-labels the unclicked samples according to the derived matching strengths, without relying on static or heuristic thresholds, thus enhancing the reliability of the generated pseudo-labels. Experimental results demonstrate that IdeFN effectively mitigates the FNS problem, achieving substantial improvements in CVR prediction accuracy.

Introduction

Conversion rate (CVR) prediction is essential in recommendation systems (Liu et al. 2021) for optimizing user engagement and personalized content delivery (Feng et al. 2024; Zhao et al. 2023; Zhang et al. 2022). Accurate CVR estimates enable systems to evaluate the likelihood of user actions (e.g., purchases or clicks) in given contexts (Huangfu et al. 2022; Sara et al. 2025), which directly informs item

*Corresponding Author.

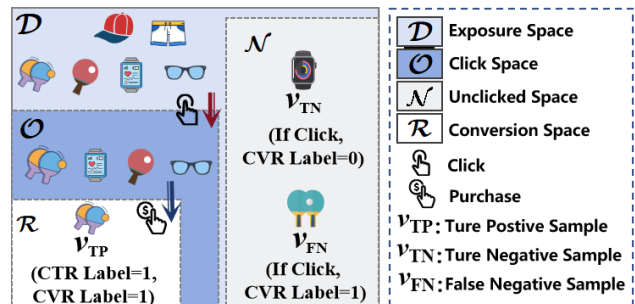


Figure 1: Illustration of key challenges in CVR prediction

ranking and enhances recommendation effectiveness (Guo et al. 2022; Chan et al. 2023; Dai et al. 2023).

In typical e-commerce scenarios, user behaviors generally follow a sequential interaction process of "exposure → click → conversion" (Wen et al. 2021). As depicted in Figure 1, the distribution of samples across the exposure space \mathcal{D} , click space \mathcal{O} , and conversion space \mathcal{R} illustrates two critical challenges inherent in this process: Sample selection bias (SSB), Data sparsity (DS). Specifically, SSB arises due to the discrepancy between training on clicked items and predicting over all exposed items, causing distribution mismatches (Dai et al. 2022; Cheng et al. 2025; Li et al. 2023). DS is caused by the scarcity of conversion events, resulting in extremely limited positive samples (Pan et al. 2025; Zhang et al. 2024). To mitigate these issues, many approaches employ multi-task learning with click-through rates (CTR) or click-through and conversion rate (CTCVR) as an auxiliary task (Ma et al. 2018b; Wang et al. 2022; Fu et al. 2024), which partially alleviates bias and sparsity.

However, since CTCVR and CVR tasks share the same label space, employing CTCVR as an auxiliary objective inevitably leads to all unclicked samples being regarded as negatives in the CVR task, which consequently gives rise to the critical issue of False negative samples (FNS). Specifically, the unclicked space \mathcal{N} comprises both true negative

samples (e.g., v_{TN} in Figure 1) and false negative samples (e.g., v_{FN} in Figure 1). The latter are not clicked due to external factors such as display order or user browsing behavior, rather than a genuine lack of interest, and may have resulted in conversions had they been clicked. Consequently, failing to distinguish these potential positives and indiscriminately labeling both v_{FN} and v_{TN} as negatives during training severely distorts the model’s characterization of user conversion preferences and ultimately degrades the accuracy of CVR prediction. Although some recent studies (Zhu et al. 2023; Su et al. 2024; Xu et al. 2022) have relaxed the assumption that all unclicked samples are negative instances, their primary focus remains on mitigating sample selection bias, and the issue of FNS in CVR prediction has received limited attention. For example, the non-click samples improved semi-supervised method (Huang et al. 2024a) predicts conversion probabilities for unclicked items as pseudo-labels to mitigate false negatives. Nevertheless, the limited integration of negative sampling into multi-task CVR frameworks and the reliance on potentially biased pseudo-labels restrict the overall effectiveness of this approach.

To address the aforementioned issues, We propose a novel multi-task framework for CVR prediction, named `IdeFN`, which identifies false negatives in the unclicked space and leverages CTR as an auxiliary task via shared-parameter learning to enhance user-behavior modeling accuracy. Moreover, `IdeFN` enables CVR prediction across the entire exposure space, primarily addressing the FNS problem while also helping to alleviate SSB and DS. Specifically, `IdeFN` consists of two main components: relaxed partial optimal transport (RPOT) module and sample relabeling mechanism (SRM). The RPOT module is a relaxed partial optimal transport algorithm that leverages a unilateral constraint and a virtual-point mechanism to learn soft matching relationships between unclicked and positive user–item pairs, thereby providing reliable signals for false-negative identification and relabeling. This design absorbs unmatched mass, suppresses spurious alignments, and maintains stability and robustness under class imbalance. SRM adaptively assigns pseudo-labels to selected unclicked samples according to their learned coupling strengths, avoiding static or heuristic thresholds and providing calibrated and reliable supervision for downstream CVR training. By integrating these two components, `IdeFN` achieves significant improvements in both prediction accuracy and model robustness.

We summarize our contributions: (1) We propose `IdeFN`, a multi-task framework for CVR prediction that explicitly targets false negatives in the unclicked space and leverages CTR as an auxiliary task, enabling accurate prediction across the entire exposure space and potentially helping to alleviate SSB and DS. (2) `IdeFN` consists of RPOT and SRM. RPOT establishes matching relationships between unclicked and positive samples to identify false negatives, and SRM adaptively generates calibrated pseudo-labels without relying on static or heuristic thresholds. (3) Extensive experiments on two datasets demonstrate the superiority of `IdeFN` over state-of-the-art baselines.

Related Work

In recent years, Multi-Task Learning (MTL) methods (Pan et al. 2019; Li et al. 2025) have been applied to CVR prediction to tackle selection bias and data sparsity (Li et al. 2025; Wang et al. 2021; Hu et al. 2025; Yi et al. 2025) by jointly optimizing multiple objectives (Li et al. 2021; Pan et al. 2022; Wei et al. 2022). Models such as ESMM (Ma et al. 2018c) and ESM² (Wang et al. 2022) treat CTR and CVR as parallel tasks with shared input features, thereby enhancing data utilization and mitigating bias. MMOE (Ma et al. 2018a) further improves MTL performance via shared and specialized expert modules. However, ESMM does not guarantee unbiased CVR estimation (Zhang et al. 2020); causal inference-based methods (e.g., Multi-IPW, Multi-DR) have been proposed to address this limitation. While these approaches alleviate selection bias in the click or exposure space, they insufficiently consider the unclick space. To address this, Zhu et al. (Zhu et al. 2023) propose a counterfactual mechanism for selection bias in the unclick space. However, existing MTL methods often exacerbate the FNS problem (Rajapakse, Yates, and de Rijke 2024; Liang et al. 2025; Huang et al. 2024b; Choudhry et al. 2022; Min et al. 2023). To mitigate this, Huang et al. (Huang et al. 2024a) introduce a semi-supervised approach that uses predicted conversion probabilities in the unclick space as pseudo-labels for CVR training. Nevertheless, reliance on potentially biased soft labels can reduce model effectiveness. Accurately identifying false negatives and developing robust CVR prediction strategies remains a key challenge.

Problem Formulation

Let $\mathcal{U} = \{u_1, \dots, u_m\}$ and $\mathcal{V} = \{v_1, \dots, v_n\}$ denote the sets of users and items, where m and n are the total numbers of users and items, respectively. $\mathcal{D} = \mathcal{U} \times \mathcal{V}$ represents the entire space consisting of all exposed user-item pairs. The click space is defined as $\mathcal{O} = \{(u, v) \mid o_{u,v} = 1, (u, v) \in \mathcal{D}\}$, where $o_{u,v}$ denotes the CTR label, with $o_{u,v} = 1$ indicating a click and $o_{u,v} = 0$ otherwise. The click matrix is represented as $\mathbf{O} \in \mathbb{R}^{m \times n}$. Similarly, the conversion space is defined as $\mathcal{R} = \{(u, v) \mid r_{u,v} = 1, (u, v) \in \mathcal{D}\}$, where $r_{u,v}$ denotes the CVR label, and $r_{u,v} = 1$ indicates indicating a conversion event (e.g., purchase), and $r_{u,v} = 0$ otherwise. The conversion matrix is denoted as $\mathbf{R} \in \mathbb{R}^{m \times n}$. In the case where all conversion labels are fully observed, the ideal loss function for CVR prediction task is:

$$\mathcal{L}_{\text{ideal}} = \frac{1}{|\mathcal{D}|} \sum_{(u,v) \in \mathcal{D}} o_{u,v} \varphi(r_{u,v}, \hat{r}_{u,v}), \quad (1)$$

where $\varphi(r_{u,v}, \hat{r}_{u,v}) = -r_{u,v} \log(\hat{r}_{u,v}) - (1 - r_{u,v}) \log(1 - \hat{r}_{u,v})$. However, as conversion labels in the unclick space $\mathcal{N} = \{(u, v) \mid o_{u,v} = 0, (u, v) \in \mathcal{D}\}$ remain unknown, CVR models are typically trained only on the click space \mathcal{O} , leading to the following naive loss:

$$\mathcal{L}_{\text{Naive}} = \frac{1}{|\mathcal{O}|} \sum_{(u,v) \in \mathcal{O}} o_{u,v} \varphi(r_{u,v}, \hat{r}_{u,v}). \quad (2)$$

It can be observed that there exists a bias between the naive CVR estimator and the ideal CVR loss. Although the CTCVR task is commonly introduced to alleviate SSB by expanding the training space from \mathcal{O} to \mathcal{D} , its label construction, defined as

$$c_{u,v} = \begin{cases} 1, & \text{if } o_{u,v} = 1 \wedge r_{u,v} = 1 \\ 0, & \text{if } o_{u,v} = 0 \text{ or } (o_{u,v} = 1 \wedge r_{u,v} = 0), \end{cases} \quad (3)$$

regards $c_{u,v} = 1$ as both a click and conversion (positive sample), and $c_{u,v} = 0$ as either no click or no conversion (negative sample). Under this label construction, optimizing CTCVR sets the CVR labels of all unclicked pairs to 0. However, unclicked samples are not necessarily true negatives, as some of them may actually correspond to future conversions. We next introduce our method for identifying these potential false negatives to improve CVR prediction across the whole exposure space.

The Proposed Method: IdeFN

To address the challenge of false negatives in the unclicked space, we propose IdeFN. As illustrated in Figure 2, the overall architecture of IdeFN consists of two main branches: the upper CTR branch, which serves as an auxiliary task to model user click behaviors, and the lower CVR branch, which is trained not only on real click-conversion samples but also on pseudo-labeled unclicked instances. Specifically, user and item features are embedded via a shared lookup table, concatenated, and then fed into a flexible backbone model shared by both CTR and CVR sub-networks for deep behavior modeling. Moreover, Figure 3 illustrates the procedure of identifying false negatives in the unclicked space via IdeFN. Specifically, RPOT constructs a relaxed partial optimal transport matrix $\tilde{\pi}^*$ to quantify the matching strength between each unclicked sample and all positive samples, thereby establishing soft coupling relationships. Based on the optimal matching matrix, SRM adaptively relabels unclicked samples to identify potential conversion instances and provides auxiliary supervision for CVR prediction.

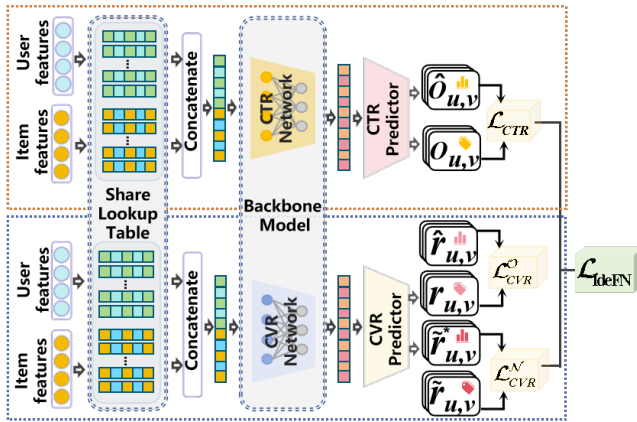


Figure 2: The overall architecture of IdeFN

RPOT: Relaxed Partial Optimal Transport

To more reliably identify false negative samples in the unclicked space, a common assumption is that these samples share similar patterns with positive samples (i.e., true conversions), allowing positive and unclicked samples to be matched for false negative discovery. Accordingly, inspired by the optimal transport (Liu et al. 2022a,b) framework based on the Kantorovich problem (Angenent, Haker, and Tannenbaum 2003), we propose the RPOT algorithm to align positive and unclicked samples and effectively identify false negatives. To start with, let $Z^P = \{z_i^P\}_{i=1}^{N_P}$ and $Z^U = \{z_j^U\}_{j=1}^{N_U}$ denote the representation sets of positive and unclicked user-item pairs, where each z_i^P and z_j^U is the concatenated representation of user and item features for the i -th positive and j -th unclicked pair, respectively. Here, N_P and N_U denote the numbers of positive and unclicked samples. Unlike the classical optimal transport assumption, which posits that the source and target domains share the same total mass, we relax this constraint to allow partial mass transport. Since unclicked samples contain both false negatives and false positives, we allow positive samples to be matched only with a subset of unclicked samples, ensuring that positives are aligned exclusively with false negatives. Accordingly, the basic RPOT can be formulated as

$$\begin{aligned} \min_{\pi \geq 0} \langle \mathcal{M}, \pi \rangle &= \min_{\pi \geq 0} \sum_{i=1}^{N_P} \sum_{j=1}^{N_U} \mathcal{M}_{ij} \pi_{ij} \\ \text{s.t. } \pi \mathbf{1}_{N_U} &= \frac{1}{N_P} \mathbf{1}_{N_P}, \quad \pi^\top \mathbf{1}_{N_P} \leq \frac{\kappa}{N_U} \end{aligned} \quad (4)$$

where $\pi \in \mathbb{R}^{N_P \times N_U}$ represents the optimal matching matrix we aim to determine. $\mathcal{M} \in \mathbb{R}^{N_P \times N_U}$ denotes the cost matrix, which is computed as $\mathcal{M}_{ij} = \|z_i^P - z_j^U\|_2^2$, with $\|\cdot\|_2$ denoting the Euclidean distance. The hyperparameter κ adjusts the proportion of false negative samples within the unclicked samples. Due to the imbalance between the numbers of positive and unclicked samples, some unclicked samples remain unmatched within the unclicked set, which are referred to as residual samples. To further ensure that the model focuses only on unclicked samples highly consistent with positives in the representation space and achieves more accurate matching, we propose a virtual point strategy that assigns unclicked samples not well-aligned with any positive to virtual points, thereby preventing samples lacking conversion potential from being incorrectly matched. Accordingly, we redefine the allocation rule for unclicked samples. The improved RPOT formulation is as follows:

$$\begin{aligned} \min_{\tilde{\pi} \geq 0} \langle \tilde{\mathcal{M}}, \tilde{\pi} \rangle &= \min_{\tilde{\pi} \geq 0} \sum_{i=1}^{N_P} \sum_{j=1}^{N_U} \tilde{\mathcal{M}}_{ij} \tilde{\pi}_{ij} \\ \text{s.t. } \Delta &= \{\tilde{\pi} \mathbf{1}_{N_U} = a = \frac{1}{N_P}, \tilde{\pi}^\top \mathbf{1}_{N_P+1} = b\} \end{aligned} \quad (5)$$

where $\tilde{\pi} \in \mathbb{R}^{(N_P+1) \times N_U}$ represents the extended matching matrix, with the additional row corresponding to the matches associated with the virtual point. Here, $b = \left(\frac{\kappa}{N_U}, 1 - \kappa \frac{N_P}{N_U}\right)$ represents the allocation ratio of unclicked samples between positive samples and the virtual point, where the latter is defined as $1 - \kappa \frac{N_P}{N_U}$, and $\tilde{\mathcal{M}} =$

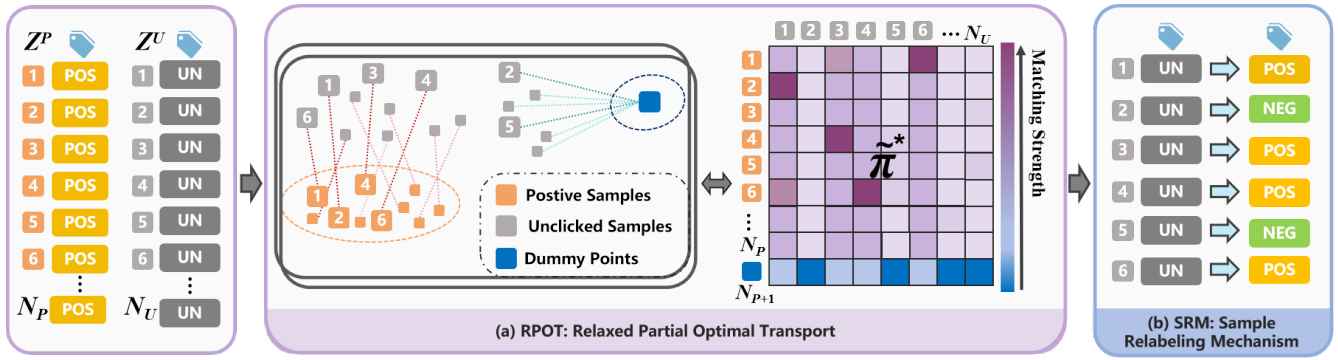


Figure 3: The basic procedure of identifying unclicked-space false negatives via IdeFN

$[\mathcal{M}, \xi \mathbf{1}_{N_U}]^\top$, denotes the extended cost matrix, where we set $\xi = 0$. As illustrated in the simplified example in Figure 3(a), the RPOT matching results are visually encoded in the matrix, where deeper box colors indicate stronger matching strengths. Specifically, unclicked samples 1, 3, 4, and 6 (gray numbers) exhibit the deepest purple boxes in rows 2, 4, 6, and 1, respectively, indicating strong associations with positive samples 2, 4, 6, and 1 (orange numbers), and are thus more likely to be considered as false negatives. Unclicked samples 2 and 5 correspond to the deepest blue boxes in the last row (dummy node), suggesting they are more likely to be true negatives.

Optimization on RPOT. Directly optimizing Eq.(5) involves solving a linear programming problem, which is computationally expensive. To address this, we introduce a KL regularization term $\mathcal{D}_{\text{KL}}(\cdot)$ to enable a more efficient optimization, as shown below:

$$\min_{\tilde{\pi}^{(\ell+1)} \in \Delta} \left\langle \tilde{\pi}^{(\ell+1)}, \tilde{\mathcal{M}} \right\rangle + \epsilon \mathcal{D}_{\text{KL}} \left(\tilde{\pi}^{(\ell+1)} \parallel \tilde{\pi}^{(\ell)} \right), \quad (6)$$

where $\tilde{\pi}^{(\ell)}$ denotes the matching matrix at the ℓ -th iteration, and ϵ denote the trade-off parameters. That is, Eq.(6) can be further rewritten as below:

$$\min_{\tilde{\pi}^{(\ell+1)} \in \Delta} \left\langle \tilde{\mathcal{M}} - \epsilon \log \tilde{\pi}^{(\ell)}, \tilde{\pi}^{(\ell+1)} \right\rangle - \epsilon \mathcal{H}(\tilde{\pi}^{(\ell+1)}), \quad (7)$$

where $\mathcal{H}(\tilde{\pi}^{(\ell+1)}) = - \langle \tilde{\pi}^{(\ell+1)}, \log(\tilde{\pi}^{(\ell+1)}) - 1 \rangle$ is the entropy regularization term. To ensure the simplicity and clarity of the subsequent formulas, we define $\tilde{\mathcal{M}}^o = \tilde{\mathcal{M}} - \epsilon \log \tilde{\pi}^{(\ell)}$. Based on Eq.(7) and using the Lagrange multiplier method we have:

$$\max_{f, g} \min_{\pi} \ell = \langle \tilde{\pi}^{(\ell+1)}, \tilde{\mathcal{M}}^o \rangle - \epsilon \mathcal{H}(\tilde{\pi}^{(\ell+1)}) - \langle f, \tilde{\pi} \mathbf{1}_{N_U} - a \rangle - \langle g, \tilde{\pi}^\top \mathbf{1}_{N_P} - b \rangle, \quad (8)$$

where, f and g are the Lagrange multipliers associated with the constraints $\tilde{\pi} \mathbf{1}_{N_U} = a$ and $\tilde{\pi}^\top \mathbf{1}_{N_P} = b$, respectively. These multipliers enforce the marginal constraints on the matching matrix $\tilde{\pi}$, ensuring that the total allocation for each unclicked sample and each positive sample adheres to the specified values of a and b , which represent the normalized proportions of positive and unclicked samples in the optimization problem. By taking the differentiation on $\tilde{\pi}_{ij}$ and

set $\frac{\partial \ell}{\partial \tilde{\pi}_{ij}} = 0$. At that time we have:

$$\tilde{\pi}_{ij}^{(\ell+1)} = \exp \left(\frac{f_i}{\epsilon} \right) \exp \left(\frac{g_j}{\epsilon} \right) \exp \left(- \frac{\tilde{\mathcal{M}}_{ij}^o}{\epsilon} \right) \geq 0. \quad (9)$$

By considering the marginal distributions provided in Eq.(10), we can iteratively solve for f and g .

$$\begin{cases} \sum_{i=1}^{N_P} \tilde{\pi}_{ij} = \exp \left(\frac{g_j}{\epsilon} \right) \sum_{i=1}^{N_P} \exp \left(\frac{f_i - \tilde{\mathcal{M}}_{ij}^o}{\epsilon} \right) = b_j \\ \sum_{j=1}^{N_U} \tilde{\pi}_{ij} = \exp \left(\frac{f_i}{\epsilon} \right) \sum_{j=1}^{N_U} \exp \left(\frac{g_j - \tilde{\mathcal{M}}_{ij}^o}{\epsilon} \right) = a_i. \end{cases} \quad (10)$$

Theorem 1. If the optimized RPOT satisfies the condition defined by the following inequality:

$$\left\langle \tilde{\mathcal{M}}, \tilde{\pi}^* \right\rangle + \eta \mathcal{D}_{\text{KL}} \left(\tilde{\pi}^* \parallel \pi^{(\ell+1)} \right) \leq \left\langle \tilde{\mathcal{M}}, \pi^{(\ell+1)} \right\rangle, \quad (11)$$

then it follows that

$$\mathcal{D}_{\text{KL}} \left(\tilde{\pi}^* \parallel \pi^{(\ell+1)} \right) \leq \frac{1}{1 + \frac{\eta}{\epsilon}} \mathcal{D}_{\text{KL}} \left(\tilde{\pi}^* \parallel \pi^{(\ell)} \right), \quad (12)$$

where η denotes a positive number.

Theorem 1 implies that the proposed optimization algorithm exhibits a linear convergence rate. A complete proof of this theorem is presented in Appendix A. Furthermore, we observe that it is unnecessary to update all variables in each iteration of the RPOT algorithm. Inspired by the Greenkhorn algorithm (Altschuler, Niles-Weed, and Rigollet 2017), we adopt an alternating optimization approach that, in each iteration, updates only the larger of the row-wise or column-wise marginal errors, thereby enhancing computational efficiency. The detailed optimization steps are presented in Algorithm 1. The time complexity of the RPOT algorithm is $O(N_P \cdot N_U)$. In each iteration, the algorithm computes the marginal errors of the transport matrix $\pi \in \mathbb{R}^{N_P \times N_U}$ along both rows and columns, and subsequently updates a single element of the scaling factor. Given a maximum iteration count ℓ_{\max} , the overall time complexity is $O(\ell_{\max} \cdot N_P \cdot N_U)$. Since N_P is usually much smaller than N_U in practice, the computational overhead remains low. Overall, through efficient optimization and accurate matching, RPOT reliably identifies false negatives, providing a solid foundation for the relabeling mechanism and enabling more accurate CVR prediction.

Algorithm 1: RPOT Algorithm

```

1: Input:  $N_P$ : Number of positive samples;  $N_U$  Number
   of unclicked samples;  $\widetilde{\mathcal{M}}^o$ : Cost matrix;  $\ell_{max}$ : Max it-
   erations;  $\epsilon, \kappa$ : Hyper parameters.
2: Initialize kernel matrix:  $K = \exp\left(-\frac{\widetilde{\mathcal{M}}^o}{\epsilon}\right)$ 
3: Initialize scaling vectors:  $u = \mathbf{1}_{N_P}, v = \mathbf{1}_{N_U}$ 
4: Initialize transport matrix:  $\widetilde{\pi} = \text{diag}(u) \cdot K \cdot \text{diag}(v)$ 
5: for  $\ell = 1$  to  $\ell_{max}$  do
6:   Compute row-wise error:  $e_{\text{row},i} = \sum_{j=1}^{N_U} \pi_{ij} - a_i$ 
7:   Compute column-wise error:  $e_{\text{col},j} = \sum_{i=1}^{N_P} \pi_{ij} - b_j$ 
8:   if  $\max(|e_{\text{row}}|) > \max(|e_{\text{col}}|)$  then
9:     Let  $i^* = \arg \max(|e_{\text{row}}|)$ 
10:    Update scaling vector:  $u_{i^*} \leftarrow \frac{a_{i^*}}{(K \cdot v)_{i^*}}$ 
11:   else
12:     Let  $j^* = \arg \max(|e_{\text{col}}|)$ 
13:    Update scaling vector:  $v_{j^*} \leftarrow \frac{b_{j^*}}{(K^T \cdot u)_{j^*}}$ 
14:   end if
15:   Update transport matrix:  $\widetilde{\pi}^* = \text{diag}(u^*) \cdot K \cdot$ 
    $\text{diag}(v^*)$ 
16: end for
17: return  $\widetilde{\pi}^* \leftarrow \widetilde{\pi}$ 

```

SRM: Sample Relabeling Mechanism

After we finish the iteration of RPOT in samples optimal transportation, we propose an adaptive relabeling mechanism (i.e., SRM) to distinguish true and false negatives and relabel the CVR labels accordingly. Specifically, for each unclicked sample j , we identify its most likely associated positive source by finding the row index i with the maximum matching mass $\widetilde{\pi}_{ij}^*$ in the optimal transport matrix $\widetilde{\pi}^*$. We define the index of the matched source as:

$$\delta_j = \arg \max_i (\widetilde{\pi}_{ij}^*), \quad (13)$$

where $\widetilde{\pi}_{ij}^*$ denotes the amount of matching mass assigned from the i -th source (either a real positive or the dummy node) to the j -th unclicked sample. Subsequently, we define the first N_P rows of the optimal matching matrix as corresponding to actual positive samples, while the (N_{P+1}) -th row is reserved for the dummy node. Based on the value of δ_j , each unclicked sample is relabeled: if $\delta_j \leq N_P$, the sample is regarded as a false negative sample; otherwise, if $\delta_j = N_{P+1}$, it is matched to the dummy node and considered true negative. The formal relabeling rule is given by:

$$\widetilde{r}_{u,v} = \begin{cases} 1, & \text{if } \delta_j \in \{1, 2, \dots, N_P\} \\ 0, & \text{if } \delta_j = N_{P+1}, \end{cases} \quad (14)$$

where $\widetilde{r}_{u,v} = 1$ and $\widetilde{r}_{u,v} = 0$ denote pseudo-positive and pseudo-negative labels, respectively. By adopting SRM, with high probability, we relabel unclicked samples 1, 3, 4, and 6 (gray) as positive (POS), indicating they are likely false negatives, while samples 2 and 5 are relabeled as negative (NEG), as shown in Figure 3(b). As depicted, samples 1, 3, 4, and 6 fall into rows 2, 4, 6, and 1, respectively, which correspond to labeled positive samples. In contrast, samples

2 and 5 fall into the (N_{P+1}) -th row, representing the dummy node, suggesting they are not aligned with any positive instance and are thus treated as true negatives.

Putting Together

To enable effective modeling over the entire exposure space, the CTR loss (\mathcal{L}_{CTR}), the CVR loss on the clicked space ($\mathcal{L}_{\text{CVR}}^{\mathcal{O}}$), and the CVR loss on the unclicked space ($\mathcal{L}_{\text{CVR}}^{\mathcal{N}}$) jointly constitute the final optimization objective of our proposed `IdEFN` in this paper:

$$\begin{aligned} \mathcal{L}_{\text{IdEFN}} &= \mathcal{L}_{\text{CTR}} + \gamma \left(\alpha \mathcal{L}_{\text{CVR}}^{\mathcal{O}} + (1 - \alpha) \mathcal{L}_{\text{CVR}}^{\mathcal{N}} \right) \\ &= \frac{1}{|\mathcal{D}|} \sum_{(u,v) \in \mathcal{D}} \varphi(o_{u,v}, \hat{o}_{u,v}) \\ &\quad + \frac{\gamma}{|\mathcal{D}|} \left(\alpha \sum_{(u,v) \in \mathcal{O}} \varphi(r_{u,v}, \hat{r}_{u,v}) \right. \\ &\quad \left. + (1 - \alpha) \sum_{(u,v) \in \mathcal{N}} \varphi(\widetilde{r}_{u,v}, \widetilde{r}_{u,v}^*) \right), \quad (15) \end{aligned}$$

where γ is a trade-off hyperparameter balancing the CTR loss and CVR-related losses, while α controls the relative importance between the CVR loss over clicked and unclicked spaces. The algorithm of `IdEFN` can be found in Appendix B Algo. 2. Overall, `IdEFN` is a collaborative multitask learning framework tailored for false negative identification in CVR prediction. With minor adaptations, it can be extended to other weakly supervised learning scenarios involving implicit feedback or noisy labels.

Experiments and Analysis

We conduct extensive experiments to answer the following questions: **Q1:** How does `IdEFN` perform compared with state-of-the-art CVR prediction methods? **Q2:** How does the performance of `IdEFN` vary with different hyperparameter settings? **Q3:** How does `IdEFN` perform when integrated into various MTL frameworks? **Q4:** What is the contribution of each component in `IdEFN` to the final performance? **Q5:** How to prove that `IdEFN` identifies false negatives via the matching mechanism in an interpretable way?

Datasets and Experimental Settings

Datasets. We conduct our experiments on two publicly available datasets **Ali-CCP** (Ma et al. 2018b) and **KuaiRand-Pure** (Gao et al. 2022), which are widely used to evaluate the CVR prediction issue. We give a detailed description of the datasets in Appendix C.1.

Baselines and Backbones¹. We compare our proposal with the following methods: **ESMM** (Ma et al. 2018b), **MMOE** (Ma et al. 2018a), **ESCM²-IPW** (Wang et al. 2022), **ESCM²-DR** (Wang et al. 2022), **DCMT** (Zhu et al. 2023), and **NISE** (Huang et al. 2024a). Besides, we choose

¹By default, we employ MLP as the backbone architecture unless specified otherwise.

| Backbone | | MLP | | | DeepFM | | | DCNV2 | | |
|---------------|------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Dataset | Method | AUC | KS | LogLoss | AUC | KS | LogLoss | AUC | KS | LogLoss |
| Ali-CCP | ESMM | 0.6313 | 0.1827 | 0.0106 | 0.6325 | 0.1819 | 0.0089 | 0.6295 | 0.1823 | 0.0099 |
| | MMoE | 0.6241 | 0.1720 | 0.0099 | 0.6216 | 0.1701 | 0.0092 | 0.6266 | 0.1765 | 0.0101 |
| | ESCM ² -IPS | 0.6439 | 0.1865 | 0.0091 | 0.6274 | 0.1967 | 0.0068 | 0.6344 | 0.2064 | 0.0085 |
| | ESCM ² -DR | 0.6110 | 0.1945 | 0.0104 | 0.6249 | 0.1350 | 0.0124 | 0.6302 | 0.2064 | 0.0144 |
| | DCMT | 0.6434 | 0.1950 | 0.0100 | <u>0.6526</u> | 0.2072 | 0.0078 | 0.6448 | 0.1688 | 0.0072 |
| | NISE | <u>0.6515</u> | <u>0.2075</u> | <u>0.0024</u> | <u>0.6523</u> | <u>0.2094</u> | <u>0.0023</u> | <u>0.6528</u> | <u>0.2117</u> | <u>0.0024</u> |
| | IdeFN | 0.6656 | 0.2238 | 0.0022 | 0.6622 | 0.2214 | 0.0021 | 0.6674 | 0.2243 | 0.0022 |
| KuaiRand-Pure | ESMM | 0.8444 | 0.5593 | 0.0703 | 0.8509 | 0.5653 | 0.0680 | 0.8436 | 0.5582 | 0.0735 |
| | MMoE | 0.8310 | 0.5447 | 0.0720 | 0.8328 | 0.5434 | 0.0688 | 0.8373 | 0.5531 | 0.0696 |
| | ESCM ² -IPS | 0.8460 | 0.5452 | <u>0.0645</u> | 0.8491 | 0.5482 | 0.0662 | 0.8449 | 0.5432 | 0.0685 |
| | ESCM ² -DR | 0.7826 | 0.4591 | <u>0.1132</u> | 0.8101 | 0.4790 | 0.1220 | 0.6442 | 0.1879 | 0.5690 |
| | DCMT | 0.8498 | 0.5535 | 0.0681 | 0.8550 | 0.5613 | <u>0.0637</u> | 0.8519 | 0.5564 | 0.0691 |
| | NISE | <u>0.8569</u> | <u>0.5584</u> | 0.0649 | <u>0.8600</u> | <u>0.5677</u> | 0.0647 | <u>0.8588</u> | <u>0.5606</u> | <u>0.0637</u> |
| | IdeFN | 0.8739 | 0.5866 | 0.0632 | 0.8779 | 0.5907 | 0.0617 | 0.8756 | 0.5883 | 0.0629 |

Table 1: Experimental results on two datasets. We bold the best result, and underline the runner-up comparison method

| Model | Method | Ali-CCP | | |
|-------|--------------------|---------------|---------------|---------------|
| | | AUC | KS | Logloss |
| ESMM | Base | 0.6313 | 0.1827 | 0.0106 |
| | +IdeFN | 0.6656 | 0.2238 | 0.0022 |
| | Improvement | 5.43% | 22.49% | 79.25% |
| MMoE | Base | 0.6241 | 0.1720 | 0.0099 |
| | +IdeFN | 0.6641 | 0.2240 | 0.0021 |
| | Improvement | 6.41% | 30.23% | 78.79% |
| Model | Method | KuaiRand-Pure | | |
| | | AUC | KS | Logloss |
| ESMM | Base | 0.8444 | 0.5593 | 0.0703 |
| | +IdeFN | 0.8739 | 0.5866 | 0.0632 |
| | Improvement | 3.49% | 4.89% | 10.09% |
| MMoE | Base | 0.8310 | 0.5447 | 0.0720 |
| | +IdeFN | 0.8672 | 0.5844 | 0.0639 |
| | Improvement | 4.36% | 7.27% | 11.17% |

Table 2: Performance of IdeFN on MTL architectures.

three backbone models, including **MLP**, **DeepFM** (Guo et al. 2017), and **DCNV2** (Wang et al. 2023) to evaluate the performance. We give a detailed introduction of baseline models and backbone models in Appendix C.2.

Evaluation Protocols and Parameter Settings. We conduct experiments on an NVIDIA GeForce RTX 4090 GPU. We adopt the Adam optimizer with early stopping to prevent overfitting. The batch size is set to 2048, the embedding dimension for each categorical feature to 16, and the learning rate to $1e-3$. All hyperparameters are selected via cross-validation: $\alpha = 0.8$, $\gamma = 30$ and $\kappa = 2e-2$. For Ali-CCP, the MLP tower is [160, 80]; for KuaiRand-Pure, [512, 256, 128, 64]. The DCNV2 model adopts a stacked structure with one cross layer, and the MMoE model is configured with 8 experts. Each experiment is run five times and average results are reported. We use the area under the ROC curve

(AUC) and LogLoss as evaluation metrics for prediction performance, and the Kolmogorov-Smirnov (KS) score to assess the model’s ability to distinguish between positive and negative samples. Higher is better for AUC and KS, whereas lower is better for LogLoss.

Performance Comparison (Q1)

The comparison results on several datasets and backbone models are shown in Table 1. From these results, we can find that: (1) IdeFN surpasses all baselines across all metrics and backbones, confirming the effectiveness of our proposal. (2) In terms of predictive accuracy, On the Ali-CCP dataset, IdeFN achieves an average relative improvement of 1.96% in AUC and a relative reduction of 8.45% in LogLoss compared to the runner-up method across MLP, DeepFM, and DCNV2. On the more challenging KuaiRand-Pure dataset, the respective gains are 2.01% (AUC) and 2.14% (LogLoss). These results confirm the effectiveness of our method across datasets and architectures. (3) In terms of discriminative ability, as measured by the KS score, IdeFN consistently outperforms the best baseline. On the Ali-CCP dataset, IdeFN achieves a mean relative improvement of 6.51%, while on KuaiRand-Pure, the improvement reaches 4.62% across the three backbones. These results indicate that IdeFN is better able to distinguish between positive and negative samples. (4) Some baseline methods, exhibit notable performance instability across different backbones. For example, ESCM²-DR experiences substantial degradation on the KuaiRand-Pure dataset when using the DCNV2 backbone, which suggests a lack of robustness and reflects an issue also reported in prior studies (Huang et al. 2024a). In summary, IdeFN’s accurate identification of false negatives enhances prediction accuracy and class discrimination, enabling reliable CVR estimation.

Parameter Study (Q2)

The impact of hyperparameters on model performance (AUC and KS) for KuaiRand-Pure is shown in Figure 4.




Figure 4: Effects of hyper-parameters α , γ and κ on KuaiRand-Pure.

We compare the results by varying $\alpha \in \{0.6, 0.7, 0.8, 0.9\}$ in (a), $\gamma \in \{10, 20, 30, 40, 50, 60, 70\}$ in (b), and $\kappa \in \{1e-3, 2e-3, 3e-3, 1e-2, 2e-2, 3e-2\}$ in (c). From the results, we observe that: (1) The best performance is achieved when $\alpha = 0.8$, $\gamma = 30$, and $\kappa = 2e-2$. (2) A small α overemphasizes the unclicked space, compromising training stability, whereas a large α leads the model to underexplore potential conversions, thereby limiting generalization. (3) A moderate γ achieves the best balance between CTR and CVR tasks. Too small a value biases training toward CTR, while too large a value lets CVR dominate, amplifying noise from pseudo-labeled unclicked samples and impairing performance. (4) The hyperparameter κ determines the proportion of unclicked samples permitted to establish matching relations with positives. An overly large value leads to excessive matching and increased label noise, whereas an insufficient value restricts the discovery of false negatives and consequently limits potential performance gains.

Evaluating IdeFN on MTL Frameworks (Q3)

We evaluate the generalization capability of the IdeFN framework across multi-task learning (MTL) architectures, including the shared-bottom model (ESMM) and the multi-expert model (MMoE). We report the results on the Ali-CCP and KuaiRand-Pure datasets in Table 2. As shown, integrating IdeFN enhances the performance of both MTL architectures across all evaluation metrics. These results underscore the robustness and adaptability of the IdeFN framework, highlighting its effectiveness in boosting predictive accuracy across diverse datasets and MTL architectures.

Ablation Study (Q4)




Figure 5: Ablation study on Ali-CCP.

To evaluate the effectiveness of each component in IdeFN, we devise the following two versions, i.e., IdeFN-T and IdeFN-R. Specifically, IdeFN-T removes the CTR auxiliary task and performs single-task CVR prediction over the exposure space. IdeFN-R removes the RPOT module and restricts CVR modeling to the clicked space. The results are shown in Figure 5. From it, we can observe that: (1) Both IdeFN-T and IdeFN-R degrade their performance compared with IdeFN. (2) The performance drop of IdeFN-T highlights the role of the CTR auxiliary task in improving user preference representation and stabilizing CVR training through auxiliary supervision. (3) IdeFN-R consistently underperforms, underscoring the need to identify false negatives in the unclicked space. Removing RPOT disables the relaxed positive matching and virtual-point strategies, hindering the discovery of high-quality false negatives and confirming RPOT’s pivotal role in extracting informative unlabeled signals and strengthening discriminative capacity.

Case Study (Q5)

We provide a detailed case study in Appendix D to further demonstrate the effectiveness of IdeFN in identifying false negatives via the matching mechanism.

Conclusion

In this paper, we propose IdeFN, a multi-task learning framework for CVR prediction that explicitly addresses the false negative sample (FNS) problem in the unclicked space. By leveraging CTR prediction as an auxiliary task, IdeFN enables accurate CVR estimation across the entire exposure space. The framework incorporates two key components: the Relaxed Partial Optimal Transport (RPOT) module, which constructs soft matching relationships between unclicked and positive samples, and the Sample Relabeling Mechanism (SRM), which adaptively generates calibrated pseudo-labels without relying on static thresholds. Extensive experiments on real-world datasets demonstrate that IdeFN significantly improves prediction accuracy and model robustness, confirming its effectiveness in mitigating the FNS issue and its applicability to practical CVR prediction scenarios.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No. 62572486), the

Natural Science Foundation of Shandong Province (No. ZR2023MF007), and by the European Union under the Horizon Europe programme (Grant No. 101135930, 101168438) and the UKRI through the UK Government's Horizon Europe Funding Guarantee (Grant No. 10119564, 10126241).

References

- Altschuler, J.; Niles-Weed, J.; and Rigollet, P. 2017. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Advances in neural information processing systems*, 30.
- Angenent, S.; Haker, S.; and Tannenbaum, A. 2003. Minimizing flows for the Monge–Kantorovich problem. *SIAM journal on mathematical analysis*, 35(1): 61–97.
- Chan, Z.; Zhang, Y.; Han, S.; Bai, Y.; Sheng, X.-R.; Lou, S.; Hu, J.; Liu, B.; Jiang, Y.; Xu, J.; et al. 2023. Capturing conversion rate fluctuation during sales promotions: A novel historical data reuse approach. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3774–3784.
- Cheng, W.; Lu, Y.; Xia, B.; Cao, J.; Xu, K.; Wen, M.; Jiang, W.; Zhang, J.; Liu, Z.; Hong, L.; et al. 2025. ChorusCVR: Chorus Supervision for Entire Space Post-Click Conversion Rate Modeling. *arXiv preprint arXiv:2502.08277*.
- Choudhry, A.; Khatri, I.; Jain, M.; and Vishwakarma, D. K. 2022. An emotion-aware multitask approach to fake news and rumor detection using transfer learning. *IEEE Transactions on Computational Social Systems*, 11(1): 588–599.
- Dai, Q.; Li, H.; Wu, P.; Dong, Z.; Zhou, X.-H.; Zhang, R.; Zhang, R.; and Sun, J. 2022. A generalized doubly robust learning framework for debiasing post-click conversion rate prediction. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 252–262.
- Dai, S.; Zhou, Y.; Xu, J.; and Wen, J.-R. 2023. Dually enhanced delayed feedback modeling for streaming conversion rate prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 390–399.
- Feng, H.; Zhang, G.; Zhang, Y.; We, Y.; and Liu, Q. 2024. EGEAN: An Exposure-Guided Embedding Alignment Network for Post-Click Conversion Estimation.
- Fu, C.; Wang, K.; Wu, J.; Chen, Y.; Huzhang, G.; Ni, Y.; Zeng, A.; and Zhou, Z. 2024. Residual Multi-Task Learner for Applied Ranking. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4974–4985.
- Gao, C.; Li, S.; Zhang, Y.; Chen, J.; Li, B.; Lei, W.; Jiang, P.; and He, X. 2022. Kuairand: an unbiased sequential recommendation dataset with randomly exposed videos. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 3953–3957.
- Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247*.
- Guo, Y.; Li, H.; Ao, X.; Lu, M.; Liu, D.; Xiao, L.; Jiang, J.; and He, Q. 2022. Calibrated Conversion Rate Prediction via Knowledge Distillation under Delayed Feedback in Online Advertising. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 3983–3987.
- Hu, W.; Sun, X.; Liu, Q.; Wu, L.; and Wang, L. 2025. Uncertainty calibration for counterfactual propensity estimation in recommendation. *IEEE Transactions on Knowledge and Data Engineering*.
- Huang, J.; Zhang, L.; Wang, J.; Jiang, S.; Huang, D.; Ding, C.; and Xu, L. 2024a. Utilizing Non-click Samples via Semi-supervised Learning for Conversion Rate Prediction. In *Proceedings of the 18th ACM Conference on Recommender Systems*, 350–359.
- Huang, J.; Zhang, L.; Wang, J.; Jiang, S.; Huang, D.; Ding, C.; and Xu, L. 2024b. Utilizing Non-click Samples via Semi-supervised Learning for Conversion Rate Prediction. In *Proceedings of the 18th ACM Conference on Recommender Systems*, 350–359.
- Huangfu, Z.; Zhang, G.-D.; Wu, Z.; Wu, Q.; Zhang, Z.; Gu, L.; Zhou, J.; and Gu, J. 2022. A multi-task learning approach for delayed feedback modeling. In *Companion Proceedings of the Web Conference 2022*, 116–120.
- Li, D.; Hu, B.; Chen, Q.; Wang, X.; Qi, Q.; Wang, L.; and Liu, H. 2021. Attentive capsule network for click-through rate and conversion rate prediction in online advertising. *Knowledge-based systems*, 211: 106522.
- Li, H.; Xiao, Y.; Zheng, C.; and Wu, P. 2023. Balancing unobserved confounding with a few unbiased ratings in de-biased recommendations. In *Proceedings of the ACM web conference 2023*, 1305–1313.
- Li, P.; Tong, X.; Wang, Y.; and Zhang, Q. 2025. Meta doubly robust: Debiasing CVR prediction via meta-learning with a small amount of unbiased data. *Knowledge-Based Systems*, 310: 112898.
- Liang, Z.; Wang, Z.; Wu, N.; Jiang, Y.; and Sun, D. 2025. A New Energy High-Impact Process Weather Classification Method Based on Sensitivity Factor Analysis and Progressive Layered Extraction. *Electronics*, 14(7): 1336.
- Liu, W.; Su, J.; Chen, C.; and Zheng, X. 2021. Leveraging distribution alignment via stein path for cross-domain cold-start recommendation. *Advances in Neural Information Processing Systems*, 34: 19223–19234.
- Liu, W.; Zheng, X.; Hu, M.; and Chen, C. 2022a. Collaborative filtering with attribution alignment for review-based non-overlapped cross domain recommendation. In *Proceedings of the ACM web conference 2022*, 1181–1190.
- Liu, W.; Zheng, X.; Su, J.; Hu, M.; Tan, Y.; and Chen, C. 2022b. Exploiting variational domain-invariant user embedding for partially overlapped cross domain recommendation. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*, 312–321.
- Ma, J.; Zhao, Z.; Yi, X.; Chen, J.; Hong, L.; and Chi, E. H. 2018a. Modeling task relationships in multi-task learning

- with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1930–1939.
- Ma, X.; Zhao, L.; Huang, G.; Wang, Z.; Hu, Z.; Zhu, X.; and Gai, K. 2018b. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 1137–1140.
- Ma, X.; Zhao, L.; Huang, G.; Wang, Z.; Hu, Z.; Zhu, X.; and Gai, K. 2018c. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 1137–1140.
- Min, C.; Lin, H.; Li, X.; Zhao, H.; Lu, J.; Yang, L.; and Xu, B. 2023. Finding hate speech with auxiliary emotion detection from self-training multi-label learning perspective. *Information Fusion*, 96: 214–223.
- Pan, J.; Mao, Y.; Ruiz, A. L.; Sun, Y.; and Flores, A. 2019. Predicting different types of conversions with multi-task learning in online advertising. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*, 2689–2697.
- Pan, X.; Li, M.; Zhang, J.; Yu, K.; Wen, H.; Wang, L.; Mao, C.; and Cao, B. 2022. MetaCVR: Conversion Rate Prediction via Meta Learning in Small-Scale Recommendation Scenarios. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, 2110–2114. Association for Computing Machinery. ISBN 978-1-4503-8732-3.
- Pan, Z.; Lou, X.; Jin, X.; Ou, C.; Liu, F.; Zeng, T.; He, C.; Liu, X.; Wei, L.; and Wang, J. 2025. Progressive Tasks Guided Multi-Source Network for Customer Lifetime Value Prediction in Online Advertising. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*, 530–538.
- Rajapakse, T. C.; Yates, A.; and de Rijke, M. 2024. Negative Sampling Techniques for Dense Passage Retrieval in a Multilingual Setting. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 575–584.
- Sara, J. D.; Pillai, J. J.; Lerman, L. O.; Lerman, A.; and Welker, K. 2025. Cardiovascular risk factors are associated with cerebrovascular reactivity in young adults. *International Journal of Cardiology*, 424: 133021.
- Su, H.; Meng, L.; Zhu, L.; Lu, K.; and Li, J. 2024. DDPO: Direct dual propensity optimization for post-click conversion rate estimation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1179–1188.
- Wang, H.; Chang, T.-W.; Liu, T.; Huang, J.; Chen, Z.; Yu, C.; Li, R.; and Chu, W. 2022. ESCM2: entire space counterfactual multi-task model for post-click conversion rate estimation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 363–372.
- Wang, W.; Feng, F.; He, X.; Zhang, H.; and Chua, T.-S. 2021. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1288–1297.
- Wang, Y.; Sun, P.; Zhang, M.; Jia, Q.; Li, J.; and Ma, S. 2023. Unbiased Delayed Feedback Label Correction for Conversion Rate Prediction. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2456–2466.
- Wei, P.; Zhang, W.; Hou, R.; Liu, J.; Liu, S.; Wang, L.; and Zheng, B. 2022. Posterior Probability Matters: Doubly-Adaptive Calibration for Neural Predictions in Online Advertising. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2645–2649.
- Wen, H.; Zhang, J.; Lv, F.; Bao, W.; Wang, T.; and Chen, Z. 2021. Hierarchically modeling micro and macro behaviors via multi-task learning for conversion rate prediction. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2187–2191.
- Xu, Z.; Wei, P.; Zhang, W.; Liu, S.; Wang, L.; and Zheng, B. 2022. Ukd: Debiasing conversion rate estimation via uncertainty-regularized knowledge distillation. In *Proceedings of the ACM Web Conference 2022*, 2078–2087.
- Yi, Q.; Tang, J.; Zhao, X.; Zeng, Y.; Song, Z.; and Wu, J. 2025. An Adaptive Entire-space Multi-scenario Multi-task Transfer Learning Model for Recommendations. *IEEE Transactions on Knowledge and Data Engineering*.
- Zhang, D.; Wu, H.; Zeng, G.; Yang, Y.; Qiu, W.; Chen, Y.; and Hu, H. 2022. CTnoCVR: A novelty auxiliary task making the lower-CTR-higher-CVR upper. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2272–2276.
- Zhang, W.; Bao, W.; Liu, X.-Y.; Yang, K.; Lin, Q.; Wen, H.; and Ramezani, R. 2020. Large-scale causal approaches to debiasing post-click conversion rate estimation with multi-task learning. In *Proceedings of The Web Conference 2020*, 2775–2781.
- Zhang, X.; Huang, C.; Zheng, K.; Su, H.; Ji, T.; Wang, W.; Qi, H.; and Li, J. 2024. Adversarial-Enhanced Causal Multi-Task Framework for Debiasing Post-Click Conversion Rate Estimation. In *Proceedings of the ACM Web Conference 2024*, 3287–3296.
- Zhao, Y.; Yan, X.; Gui, X.; Han, S.; Sheng, X.-R.; Yu, G.; Chen, J.; Xu, Z.; and Zheng, B. 2023. Entire Space Cascade Delayed Feedback Modeling for Effective Conversion Rate Prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 4981–4987.
- Zhu, F.; Zhong, M.; Yang, X.; Li, L.; Yu, L.; Zhang, T.; Zhou, J.; Chen, C.; Wu, F.; Liu, G.; et al. 2023. DCMT: A Direct Entire-Space Causal Multi-Task Framework for Post-Click Conversion Estimation. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 3113–3125.