PARS: Partial-Label-Learning-Inspired Recommender Systems

Shanshan Ye, Kezhi Lu, Guangquan Zhang, Jie Lu
Australian Artificial Intelligence Institute (AAII),
University of Technology Sydney, Ultimo, NSW 2007, Australia
{shanshan.ye@student, kezhi.lu-1@student, guangquan.zhang, jie.lu}@uts.edu.au

Abstract

Recommender systems are widely required and deployed to address real-world problems. In this paper, we study a new yet challenging real-world setting for recommender systems, where only user browsing histories are available without any explicit feedback. No item acquisition information, e.g., purchasing or rating, is given. By assuming that user browsing sequences are likely to contain the items to acquire, we draw an analogy to the setting of partial label learning in weakly supervised learning. This enables us to train reliable recommender systems only using browsing histories. We term the proposed method as Partial Acquisition Recommender System (PARS). Empirical results on real-world benchmark datasets show the effectiveness of the proposed method. Surprisingly, we also show that the proposed method even surpasses some baselines using item acquisition information.

Introduction

Recommender systems (RS) have become essential to our modern lives, significantly reshaping how individuals acquire products, services, and content. Fundamentally, these systems act as information filtering tools, designed to provide personalised products or services, helping to address the growing problem of information overload (Resnick and Varian 1997). RS have various applications (Lu et al. 2015), e.g., from prominent digital commerce such as Amazon and Alibaba to powerful media streaming platforms like Netflix and YouTube (Roy and Dutta 2022). Their pervasive impact highlights their substantial market value and critical role in improving user engagement, increasing sales, and enhancing satisfaction, establishing them as a foundational element of modern e-commerce and digital content delivery.

To achieve satisfactory application performances, RS need to capture user interests precisely (Claypool et al. 2001; Lu et al. 2024, 2025). Traditionally, RS rely heavily on explicit acquisition feedback (Liang, Pan, and Ming 2021), such as preferences, ratings, purchasing information, etc., to build effective recommendation models. However, in real-world settings, due to data accessibility (Huang et al. 2024b) or privacy factors (Dwork 2006), there are plenty of cases where explicit acquisition feedback is unavailable (Núñez-Valdez et al. 2018). For example, users may browse product pages, read articles, or watch video clips online without explicitly indicating a preference, leaving a rating, or making a purchase. For these situations where only browsing histories are available without any explicit acquisition information, it is demanding and important to build effective RS.

Although the aim of building RS based solely on browsing histories is compelling and appealing, it is infeasible to achieve it without any assumptions or constraints. For example, if users randomly browse products that they are not interested in, it is impossible to learn any useful item acquisition information from such browsing data. To tackle this issue, we make a crucial and realistic assumption: users are more likely to browse products that interest them, i.e., the majority of browsing histories contain items that the users would like to acquire. We will discuss the assumption later to show it is weak. If this hidden item acquisition information can be effectively extracted from the browsing histories, it becomes feasible to build accurate and robust RS without relying on explicit feedback.

To understand why the latent item acquisition information can be extracted, we draw an analogy of the proposed setting to a weakly supervised learning setting termed as partial label learning (Xu, Lv, and Geng 2019). In partial label learning, each instance has been assigned a set consisting of multiple labels while the set contains the true label. The goal is to learn a classifier that can predict true labels and generalize to unseen instances. It has been proven that effective learning is feasible for this setting (Cour, Sapp, and Taskar 2011). Later studies have extended the setting where the majority of label sets contain the true labels (Xu et al. 2023). It is intuitive that the weak supervision information, i.e., the majority of label sets contain true labels, can be exploited to build a reliable classifier. In our setting, we have an analogy that the majority of user browsing histories contain the items to acquire, making it feasible to extract item acquisition information to build reliable RS.

In this paper, we propose Partial Acquisition Recommender Systems (PARS), a novel framework that, for the first time, enables reliable RS only exploiting browsing histories. The key innovation of PARS lies in its ability to learn conversion probabilities (or item purchase probabilities) from browsing to transaction without directly giving supervision of item acquisition. As shown in Figure 1, PARS introduces a partial label learning module combined with a

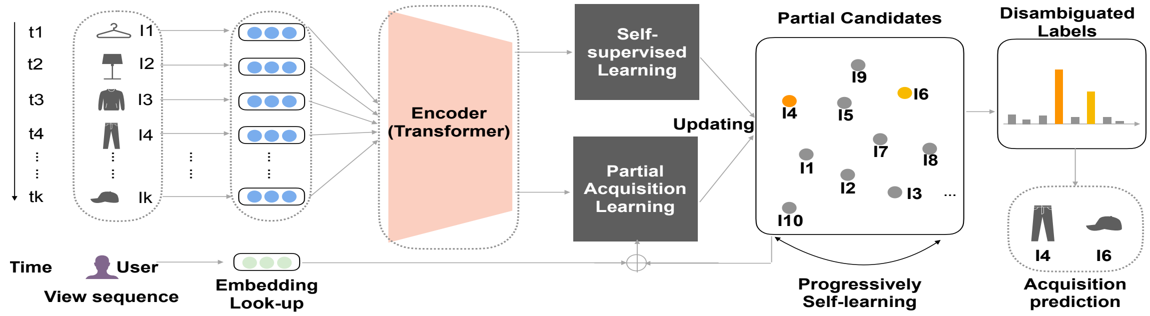

Figure 1: Illustration of the proposed PARS framework. The self-supervised learning part continuously updates the item embeddings. The partial label learning part is used to update the pseudo conversion probabilities gradually.

self-learning mechanism that iteratively refines conversion probability estimates for items in user browsing sequences. Specifically, a masked self-learning part is used in the self-supervised learning part to capture temporal dynamics and contextual patterns in user browsing sequences. The partial label learning part ensures that initial noisy estimates are gradually refined as the model exploits the weakly item acquisition information that the majority of user browsing sequences contain the items to acquire. This approach represents a paradigm shift from supervised to unsupervised RS, making PARS particularly suitable for real-world scenarios where conversion feedback is unavailable.

Empirical results on popular real-world benchmark datasets consistently show the effectiveness of PARS. Since PARS is the first methodology to enable building RS only using browsing histories, there is no straightforward baseline. A related topic is conversion rate (CVR) prediction (Wen et al. 2019), which aims to estimate the probability that a user will complete a purchase after being exposed to an item. However, CVR needs the supervision of direct item acquisition labels in the training data. We show that PARS, without using any direct item acquisition information, will outperform CVR methods which use item acquisition information during the training phase.

Our work also has interesting implications. Firstly, it tackles a specific form of cold-start problem, where RS struggle to make predictions for new users with no acquisition data but do have early browsing histories (Rajendran and Sundarraj 2021). By utilizing only browsing histories, our method can produce meaningful recommendations. Secondly, our approach introduces a novel privacy-preserving mechanism. Instead of collecting and storing sensitive item acquisition data such as purchase records or explicit ratings, it relies solely on user browsing histories, which are generally less privacy-intrusive and can still yield strong recommendation performance. As a result, this capability fosters greater user trust and expands their applicability to a wider range of real-world contexts where explicit item purchase information is either limited by policy or considered ethically inappropriate (Ye, Lu, and Zhang 2025).

In summary, our main contributions are as follows: (1) we formalize the problem of building RS using only browsing histories; (2) we propose a principled framework to recover acquisition information under realistic assumptions; and (3) we empirically validate our approach through comprehensive experiments, showing its robustness and real-world applicability.

Related Work

In this section, we review relevant literature in three key areas: conversion rate prediction, self-supervised learning for sequential recommendation (SR), and partial label learning.

Conversion Rate Prediction Conversion rate (CVR) prediction (Wen et al. 2019) aims to estimate the probability that a user will complete a purchase after being exposed to an item. CVR faces two fundamental challenges: sample selection bias, e.g., (Ma et al. 2018b,a; Wen et al. 2020; Wang et al. 2022; Su et al. 2024; Huang et al. 2024a), and data sparsity, e.g., (Guo et al. 2017; Zhu et al. 2023; Ouyang et al. 2023), as conversion labels are only observed for purchased items.

Despite their success, all existing CVR methods share a fundamental requirement: they require explicit conversion labels during training. This requirement is frequently violated in practice due to delayed conversions, cross-device purchases, and privacy constraints, etc. Our work challenges this requirement by proposing a method that can learn purchase patterns without any explicit conversion labels.

Self-Supervised Learning for Sequential Recommender Systems Our study involves user browsing sequences (Ye and Lu 2023). Therefore it is closely related to sequential RS (Wang et al. 2019), which aim to predict users' future interactions based on their historical behavior sequences. Recent advances have shown that self-supervised learning (Gui et al. 2024) can significantly improve representation learning without requiring explicit item acquisition labels.

Inspired by masked language modeling in NLP, masked sequence modeling has been proposed to SR (Kang and McAuley 2018; Sun et al. 2019). Another line of work employs contrastive learning to enhance sequence representations (Xie et al. 2022; Liu et al. 2021).

Since our setting does not have any access to explicit item acquisition information, we employ masked sequence modeling to help learn robust item embeddings.

Partial Label Learning Partial label learning (PLL) (Xu, Lv, and Geng 2019) is a specific weakly supervised classification problem (Zhou 2018; Ye and Lu 2024). In PLL, each training instance is associated with a set of candidate labels, one of which is correct. This paradigm naturally handles label ambiguity in real-world applications.

Early PLL approaches can be categorized into two strategies: averaging-based methods (Cour, Sapp, and Taskar 2011) that treat all candidate labels equally, and disambiguation-based methods that iteratively refine label confidences (Zhang, Zhou, and Liu 2016). Theoretical analyses have been provided to show the learnability, consistency, and generalization of PLL algorithms (Cour, Sapp, and Taskar 2011; Lv et al. 2020; Feng et al. 2020; Xu et al. 2021).

Existing PLL methods are primarily designed for multi-class classification with balanced datasets. We are the first to introduce it to the community of RS.

Proposed Method

In this section, we first setup the research problem. Then, we discuss the assumption enabling the idea of partial label learning (PLL) to extract item acquisition information. To learn user and item embeddings to estimate item acquisition probabilities, we design a transformer-based framework and employ masked language modeling and PLL to help learn semantic embeddings in an unsupervised manner.

Problem Formulation Let $\mathcal{U} = \{u_1, u_2, ..., u_n\}$ denote the set of n users, and $\mathcal{I} = \{i_1, i_2, ..., i_m\}$ denote the set of m items. Each user $u \in \mathcal{U}$ is associated with a session consisting of a sequence of interacted items $\mathcal{S}_u = \left\{i_1^u, i_2^u, ..., i_{|\mathcal{S}_u|}^u\right\}$, where $i_t^u \in \mathcal{I}$ represents the t-th item in user u's interaction sequence. Let $L = \max\left\{|\mathcal{S}_u| \;\middle|\; u \in \mathcal{U}\right\}$ representing the maximum sequence length. We padd all the sequence to the length of L, i.e., $\mathcal{S}_u = \{i_1^u, i_2^u, ..., i_L^u\}, \forall u \in \mathcal{U}$.

We aim to learn which items users are likely to purchase based only on the observed interaction sequences $\{\mathcal{S}_u | u \in \mathcal{U}\}$. This setting is quite realistic in many e-commerce scenarios where purchase signals are delayed, missing, or privacy-protected. Given a user u with interaction sequence $\mathcal{S}_u$, our goal is to predict an acquisition vector $\hat{\mathbf{p}}_u \in \mathbb{R}^L$ where $\hat{\mathbf{p}}_{u,i}$ represents the probability that user u will purchase item $i, i \in \{1, 2, \ldots, L\}$. The learning problem can be formulated as

$$f_\theta : \mathcal{S}_u \mapsto \hat{\mathbf{p}}_u \tag{1}$$

where θ represents the parameters of model f. We term the problem as Partial Acquisition Recommender Systems (PARS) learning.

An Analogy to Partial Label Learning We introduce the setting of PLL and then draw an analogy of PARS to PLL.

We misuse notation slightly to set up PLL for easy comparison. Specifically, let $\mathcal{X}$ be the instance space and $\mathcal{Y}$ the label set of multiple classes. Given training data $\{(\mathbf{x}_i, \mathcal{S}_i) | 1 \le i \le n\}$, where $\mathbf{x}_i$ represents an instance, $\mathcal{S}_i$ a set of L possible labels, n the size of instances, PLL aims to learn a multi-class classifier $f : \mathcal{X} \to \mathcal{Y}$. It has been proven that if the majority of the label set $\mathcal{S}$ contains true labels, a reliable classifier f can be learned (Xu et al. 2023).

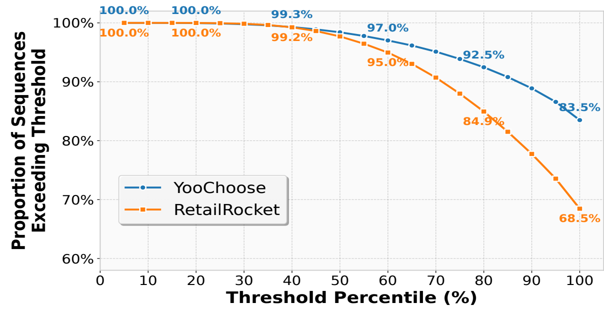

Figure 2: The vertical axis shows the percentage of sequences that contain at least one either positive or pseudo positive acquisition label. The horizontal axis represents different ways to calculate the threshold for pseudo positvie acquisition labels. $x\%$ means the threshold equals the average value of the smallest $x\%$ of the predicted scores of the items with true positive acquisition labels. Note 100% means the threshold equals the average value of all the predicted scores of the items with true positive acquisition labels.

We can see that PARS and PLL have very similar settings, where users in PARS correspond to instances in PLL; an item sequence that contains items the user may purchase corresponds to a label set that contains the possible labels for an instance; both settings have to exploit the underlying true acquisition/label information from the interaction sequence/label set. However, in PLL, it is required that the majority of label sets should at least contain one true label; while in PARS, it is not clear if the majority of sequences will contain at least one item purchased or not, because PARS study real-world applications where data distribution cannot be controlled.

If **we assume that the majority of sequences contain at least one item purchased**, we can adapt the methodologies designed for PLL for PARS. It looks like the assumption is strong. However, in the next subsection, we will show that the popular benchmark datasets in RS used in the paper implicitly satisfy the assumption.

The Assumption is Weak In this paper, we study the three popular datasets in RS, i.e., YooChoose[1], RetailRocket[2], and Taobao[3]. More details about the dataset are presented in the supplementary material.

All the three dataset have the ground-truth item acquisition labels. Note that for PARS, we do not use the labels to train the model. We found that Taobao has 75.45% of the interaction sequences containing at least one item marked as purchased. However, YooChoose and RetailRocket have less than 50%, i.e., YooChoose has 20.98% and RetailRocket has 5.57%. It seems the assumption does not hold that the majority of sequences contain at least one item purchased. However, we argue that in real-world scenarios, the browsed item sequences are likely to contain the item the user would like to purchase, even though the transactions have not been

[1]https://www.kaggle.com/chadgostopp/recsys-challenge-2015
[2]https://www.kaggle.com/retailrocket/ecommerce-dataset
[3]https://tianchi.aliyun.com/dataset/649?lang=en-us

made, e.g., the users may purchase the items on other platforms or at a later time.

To verify the above argument, we have checked the predicted item acquisition probabilities $\hat{\mathbf{p}}_u$ for the training data and found that many items have predicted acquisition probabilities higher than those with positive ground-truth item acquisition labels. We can assign pseudo positive acquisition labels to the items. As shown in Figure 2, we can see that the vast majority of items have at least one either positive or pseudo positive acquisition label, making our above assumption implicitly hold true. Therefore, the methodologies designed for PLL can be borrowed to address PARS. In the following, we will show how to learn user and item embeddings and how to use the PLL philosophy for PARS.

Transformer-based User Embedding To obtain user embeddings capturing item and sequential information, we use the following item and position embedding

$$\mathbf{e}_t = \mathbf{Z}_{\text{item}}(i_t^u) + \mathbf{Z}_{\text{pos}}(t) \tag{2}$$

where $\mathbf{Z}_{\text{item}} \in \mathbb{R}^{d\times m}$ is the item embedding matrix (with an additional token embedding for masking), $\mathbf{Z}_{\text{pos}} \in \mathbb{R}^{d\times L}$ is the positional embedding matrix, and d is the dimension of the embedding, $\mathbf{Z}(x)$ represents the x-th column of $\mathbf{Z}$.

We further employ a multi-layer Transformer encoder to enhance the expressive power of the embeddings, i.e.,

$$\mathbf{H}^{(l)} = \text{TransformerLayer}^{(l)}(\mathbf{H}^{(l-1)}), \quad l = 1, ..., N \tag{3}$$

where $\mathbf{H}^{(0)} = [\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_L]$, N is the number of Transformer layers, and each layer consists of multi-head self-attention and feed-forward networks with residual connections and layer normalization.

To obtain a global representation of the user's sequence, we employ the average pooling over the non-padding part of $\mathbf{H}^{(N)}$, i.e.,

$$\mathbf{g}_u = \frac{1}{\sum_{t=1}^{L}(1-m_t)} \sum_{t=1}^{L} (1-m_t)\cdot \mathbf{H}_t^{(N)} \tag{4}$$

where $\mathbf{H}_t^{(N)}$ is the t-th hidden state from the final Transformer layer (the t-th column of $\mathbf{H}^{(N)}$), and $m_t \in \{0,1\}$ is the indicator for padding.

We then project the global representation to obtain the final user embedding for user u, i.e.,

$$\mathbf{z}_u = \text{LayerNorm}(\mathbf{W}_g \mathbf{g}_u + \mathbf{b}_g) \tag{5}$$

where $\mathbf{W}_g \in \mathbb{R}^{d\times d}$ is a projection matrix, $\mathbf{b}_g$ is a bias vector, and LayerNorm$(\cdot)$ represents a normalization layer.

Masked Language Modeling for Semantic Embeddings To learn semantic embeddings, we employ the masked language modeling (MLM) (Devlin et al. 2019). During training, we randomly mask a subset of items in each sequence and train the model to predict the masked items.

For a sequence $\mathcal{S}_u$, we randomly select positions $\mathcal{M} \subset \{1, 2, ..., L\}$ to mask with probability p_{mask}. The masked sequence $\tilde{\mathcal{S}}_u$ is created by replacing items at positions in $\mathcal{M}$ with a special [MASK] token.

MLM learns semantic embedding by minimizing the distance between the predicted masked items and the ground truth. The objective is computed as

$$\mathcal{L}_{\text{MLM}} = -\frac{1}{|\mathcal{M}|}\sum_{t\in\mathcal{M}} \log p\left(i_t^u \mid \tilde{\mathbf{H}}_t^{(N)}\right) \tag{6}$$

where $\tilde{\mathbf{H}}_t^{(N)}$ is the hidden representation at position t from the masked sequence, and

$$p\left(i|\tilde{\mathbf{H}}_t^{(N)}\right) = \frac{\exp\left(\mathbf{w}_i^\top \tilde{\mathbf{H}}_t^{(N)}\right)}{\sum_{j=1}^{m}\exp\left(\mathbf{w}_j^\top \tilde{\mathbf{H}}_t^{(N)}\right)} \tag{7}$$

with $\mathbf{w}_i$ being the i-th row of the MLM prediction head $\mathbf{W}_{\text{MLM}} \in \mathbb{R}^{m\times d}$.

Partial Label Learning for Purchase Prediction Since we lack explicit item acquisition labels during the model training phase, we use the idea of partial label learning to predict the item acquisition probability.

We maintain a partial label matrix $\bar{\mathbf{P}} = [\bar{\mathbf{p}}_1, \bar{\mathbf{p}}_2, \ldots, \bar{\mathbf{p}}_n] \in [0,1]^{m\times n}$ where $\bar{\mathbf{p}}_{u,i}$ indicates the probability that user u will purchase item i. We initialize partial labels $\bar{\mathbf{P}}$ by uniformly distributing probability mass across all items observed in each user's interaction sequence. Specifically, for each user u with interaction sequence $\mathcal{S}_u = \{i_1^u, i_2^u, \ldots, i_L^u\}$, we set

$$\bar{\mathbf{p}}_{u,i}^{(0)} = \begin{cases} \frac{1}{|\mathcal{U}_u|}, & \text{if } i \in \mathcal{U}_u \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

where $\mathcal{U}_u = \{i \mid i \in \mathcal{S}_u\}$ represents the set of unique items in user u's interaction sequence.

Given a partial label matrix $\bar{\mathbf{P}}$, PLL uses the following objective computed using cross-entropy between the model predictions $\hat{\mathbf{p}}_u$ and the partial labels $\bar{\mathbf{P}}$ to update $\hat{\mathbf{p}}_u$, i.e.,

$$\mathcal{L}_{\text{PLL}} = -\frac{1}{n}\sum_{u=1}^{n}\sum_{i=1}^{m} \bar{\mathbf{p}}_{u,i}\cdot \log(\hat{\mathbf{p}}_{u,i} + \epsilon) \tag{9}$$

where ϵ is a small constant (i.e., 10^{-10}) for numerical stability, $\hat{\mathbf{p}}_u^{(t)} = \text{softmax}\left(\mathbf{Z}_{\text{item}}^\top \mathbf{z}_u\right)$, $\mathbf{Z}_{\text{item}} \in \mathbb{R}^{d\times m}$ contains the item embeddings of all items.

PLL alternatively updates the partial labels $\bar{\mathbf{P}}$ and the model predictions $\hat{\mathbf{p}}_u$. Given the current model predictions $\hat{\mathbf{p}}_u^{(t)}$. The partial labels $\bar{\mathbf{P}}$ are then refined through a two-step process. In the first step,

$$\tilde{\mathbf{p}}_{u,i}^{(t+1)} = \begin{cases} \hat{\mathbf{p}}_{u,i}^{(t)} & \text{if } \bar{\mathbf{p}}_{u,i}^{(t)} > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

Note that $\tilde{\mathbf{p}}_u^{(t+1)}$ is further normalized to maintain a valid probability distribution.

In the second step, to ensure stable training, we employ the momentum method, i.e.,

$$\bar{\mathbf{p}}_{u,i}^{(t+1)} = \alpha\cdot\tilde{\mathbf{p}}_{u,i}^{(t+1)} + (1-\alpha)\cdot\bar{\mathbf{p}}_{u,i}^{(t)} \tag{11}$$

where $\alpha \in [0,1]$ is the momentum coefficient that controls the trade-off between historical and current predictions.

Algorithm 1: Partial Acquisition Recommender Systems

Require: User sequences $\{\mathcal{S}_u\}$, number of epochs T, learning rate η, MLM weight λ
Ensure: Trained model parameters θ
Initialize model parameters θ randomly
Initialize partial labels $\mathbf{Y}^{(0)}$
for epoch $= 1$ to T **do**
  for each batch of sequences **do**
    // Masked Language Modeling
    Create masked sequences
    Compute sequence representations $\mathbf{H}$
    Compute $\mathcal{L}_{\text{MLM}}$ using masked positions
    // Partial Label Learning
    Compute user embedding $\mathbf{z}_u$
    Compute prediction scores $\hat{\mathbf{p}}_u$
    Update partial labels $\bar{\mathbf{P}}$
    Compute $\mathcal{L}_{\text{PLL}}$ using updated partial labels
    // Parameter Update
    $\mathcal{L} = \mathcal{L}_{\text{PLL}} + \lambda \cdot \mathcal{L}_{\text{MLM}}$
    $\theta = \theta - \eta \cdot \nabla_\theta \mathcal{L}$
  end for
end for
return θ

Partial Acquisition Recommender Systems Optimizing the objectives of MLM and PLL will enable the Transformer-based user embedding framework to learn meaningful user and item embeddings. Specifically, MLM will help capture sequential information. With the assumption that the majority of sequences contain at least one item to purchase, PLL will help extract reliable item acquisition information, i.e., $\hat{\mathbf{p}}_{u,i}$ will be large if the user u is likely to purchase the item i and $\hat{\mathbf{p}}_{u,i}$ will be small in the other way. We therefore build the objective of Partial Acquisition Recommender Systems (PARS) as follows:

$$\mathcal{L} = \mathcal{L}_{\text{PLL}} + \lambda \cdot \mathcal{L}_{\text{MLM}} \tag{12}$$

where λ is a hyperparameter controlling the weight of the MLM objective. The complete training procedure is summarized in Algorithm 1.

At the inference phase, given a user's interaction sequence $\mathcal{S}_{u_{\text{new}}}$, we can use the trained Transformer encoder to obtain the user embedding $\mathbf{z}_{u_{\text{new}}}$. We then can calculate $\hat{\mathbf{p}}_{u_{\text{new}}} = \mathbf{z}_{u_{\text{new}}}^\top \mathbf{Z}_{\text{item}}$. PARS ranks the items according to $\hat{\mathbf{p}}_{u_{\text{new}}}$ to generate purchase recommendations.

The key advantage of our approach is that it can learn meaningful purchase patterns through self-supervision and partial label learning, without requiring explicit item acquisition labels during training. This makes it applicable to a wide range of real-world scenarios where obtaining complete item acquisition labels is challenging or impossible.

Experiments

Datasets To comprehensively evaluate our proposed method, we conduct extensive experiments on three large-scale public datasets, i.e., YooChoose, RetailRocket, and Taobao, collected from real-world e-commerce platforms. These datasets represent diverse user behaviors and purchase patterns across different platforms and regions.

Dataset	#Users	#Items	#Interactions	Avg. Length
YooChoose	417,370	39,299	6,197,045	14.85
RetailRocket	116,679	112,574	620,338	5.32
Taobao	213,374	68,777	3,563,298	16.70

Table 1: Statistics of the experimental datasets.

Table 1 summarizes the statistics of the three datasets after preprocessing (e.g., removing items with fewer than 5 interactions; more details can be found in the supplementary materials). The datasets vary significantly in scale and characteristics, with YooChoose having the most sessions but fewer items, RetailRocket being more sparse with shorter sequences, and Taobao containing rich user behavior patterns with longer interaction sequences.

Evaluation Metrics We adopt a comprehensive set of evaluation metrics, i.e., AUC, HR@k (Hit Rate@k), NDCG@k (Normalized Discounted Cumulative Gain@k), precision@k, recall@k, F1@k, to assess the performance of our method from multiple perspectives. The definitions are postponed in the supplementary materials. We evaluate all metrics with $k \in \{1, 3, 5\}$ to assess performance at different recommendation list lengths.

Baseline Methods Note that PARS is the first methodology to enable building recommender systems only using browsing histories; there is no straightforward baseline. We compare PARS against ten state-of-the-art CVR baseline methods that using item acquisition labels.

- **DeepFM** (Guo et al. 2017): Combines factorization machines with deep neural networks to capture both low-order and high-order feature interactions.
- **ESMM** (Ma et al. 2018b): Entire Space Multi-Task Model addresses sample selection bias by modeling CVR over entire exposure space rather than clicked samples.
- **MMoE** (Ma et al. 2018a): Multi-gate Mixture-of-Experts learns task-specific and shared representations through a gating mechanism for multi-task learning.
- **BERT4Rec** (Sun et al. 2019): Bidirectional Encoders from Transformer for SR that uses bidirectional self-attention to model user behavior sequences and employs Cloze task for pre-training to capture contextual information from both directions.
- **ESM^2** (Wen et al. 2020): Entire Space Multi-task Modeling with field-level calibration to handle task conflicts and improve CVR prediction.
- **ESCM^2** (Wang et al. 2022): Entire Space Counterfactual Multi-task Modeling that employs counterfactual risk minimization to address selection bias and data sparsity.
- **DCMT** (Zhu et al. 2023): Disentangled Causal Multi-Task learning that separately models CVR in click and non-click spaces to reduce gradient conflicts.
- **CL4CVR** (Ouyang et al. 2023): Contrastive Learning for CVR prediction that leverages self-supervised learning to improve representation learning.

Dataset	Method	AUC	HR@K			NDCG@K		Precision@K		Recall@K			F1@K		
			@1	@3	@5	@3	@5	@3	@5	@1	@3	@5	@1	@3	@5
YooChoose	DeepFM	<u>0.6964</u>	<u>0.5379</u>	<u>0.8334</u>	0.9085	<u>0.5438</u>	<u>0.6053</u>	<u>0.4768</u>	<u>0.4140</u>	<u>0.1693</u>	<u>0.4381</u>	<u>0.6299</u>	<u>0.2407</u>	<u>0.4193</u>	<u>0.4599</u>
	ESMM	0.5970	0.5016	0.8164	0.9017	0.5226	0.5881	0.4605	0.4045	0.1611	0.4247	0.6181	0.2273	0.4060	0.4501
	MMOE	0.6724	0.4870	0.8108	0.8994	0.5132	0.5767	0.4555	0.3982	0.1547	0.4188	0.6083	0.2191	0.4008	0.4429
	BERT4Rec	0.5786	0.4066	0.7332	0.8741	0.4325	0.5049	0.3875	0.3690	0.1289	0.3590	0.5641	0.1822	0.3412	0.4104
	ESM2	0.6945	0.5070	0.8280	0.9065	0.5334	0.5968	0.4721	0.4119	0.1616	0.4359	0.6282	0.2284	0.4163	0.4583
	ESCM^2-DR	0.6327	0.4926	0.8199	0.9054	0.5193	0.5867	0.4610	0.4067	0.1557	0.4238	0.6214	0.2210	0.4058	0.4529
	ESCM^2-IPS	0.6621	0.5063	0.8282	0.9096	0.5287	0.5943	0.4670	0.4094	0.1598	0.4325	0.6268	0.2271	0.4125	0.4562
	DCMT	0.6687	0.4944	0.8211	0.9023	0.5226	0.5866	0.4640	0.4060	0.1556	0.4288	0.6194	0.2213	0.4094	0.4519
	CL4CVR	0.6635	0.5006	0.8278	0.9102	0.5269	0.5935	0.4653	0.4091	0.1606	0.4341	0.6279	0.2270	0.4122	0.4564
	DDPO	0.6838	0.5026	0.8274	0.9063	0.5244	0.5890	0.4617	0.4043	0.1614	0.4289	0.6211	0.2277	0.4082	0.4511
	NISE	0.5946	0.5012	0.8256	<u>0.9116</u>	0.5260	0.5909	0.4616	0.4045	0.1624	0.4329	0.6235	0.2283	0.4095	0.4516
	EVI	0.6597	0.4888	0.8177	0.9044	0.5177	0.5826	0.4595	0.4028	0.1545	0.4241	0.6170	0.2191	0.4052	0.4489
	PARS (Ours)	**0.9398**	**0.6173**	**0.8530**	**0.9204**	**0.5903**	**0.6285**	**0.4868**	**0.4170**	**0.2278**	**0.4761**	**0.6393**	**0.3066**	**0.4395**	**0.4645**
RetailRocket	DeepFM	0.5061	0.2965	0.6852	0.8499	0.4702	0.5512	0.2884	0.2344	0.1980	0.5934	0.7847	0.2231	0.3647	0.3407
	ESMM	0.5188	0.3207	0.7146	0.8707	0.4978	0.5769	0.2987	0.2385	0.2242	0.6184	0.8039	0.2492	0.3791	0.3477
	MMOE	0.5166	0.3063	0.7023	0.8503	0.4813	0.5584	0.2899	0.2336	0.2121	0.6048	0.7846	0.2365	0.3691	0.3402
	BERT4Rec	0.5103	0.3156	0.7244	0.8664	0.4996	0.5732	0.2994	0.2363	0.2198	0.6289	0.8011	0.2452	0.3828	0.3454
	ESM2	<u>0.5487</u>	0.3650	<u>0.7376</u>	0.8792	<u>0.5311</u>	<u>0.6040</u>	<u>0.3085</u>	<u>0.2427</u>	0.2612	<u>0.6453</u>	<u>0.8160</u>	0.2884	<u>0.3929</u>	<u>0.3532</u>
	ESCM^2-DR	0.5156	0.3207	0.7027	0.8562	0.4936	0.5723	0.2934	0.2340	0.2193	0.6077	0.7893	0.2455	0.3717	0.3410
	ESCM^2-IPS	0.5217	0.3041	0.6891	0.8537	0.4762	0.5581	0.2874	0.2350	0.2084	0.5961	0.7873	0.2331	0.3647	0.3418
	DCMT	0.5246	0.3271	0.6976	0.8677	0.4917	0.5753	0.2921	0.2373	0.2290	0.6047	0.8011	0.2549	0.3707	0.3459
	CL4CVR	0.5179	0.3131	0.6937	0.8541	0.4850	0.5659	0.2914	0.2352	0.2203	0.6018	0.7909	0.2444	0.3693	0.3427
	DDPO	0.5468	<u>0.3764</u>	0.7350	<u>0.8796</u>	0.5357	0.6078	0.3061	0.2393	<u>0.2777</u>	0.6425	0.8127	<u>0.3037</u>	0.3907	0.3498
	NISE	0.5265	0.3296	0.7095	0.8673	0.4988	0.5789	0.2956	0.2390	0.2287	0.6169	0.8043	0.2545	0.3760	0.3482
	EVI	0.5079	0.3271	0.7091	0.8562	0.4955	0.5723	0.2934	0.2354	0.2306	0.6108	0.7910	0.2556	0.3729	0.3426
	PARS (Ours)	**0.9718**	**0.7835**	**0.9145**	**0.9507**	**0.8076**	**0.8379**	**0.3902**	**0.2687**	**0.6194**	**0.8286**	**0.9030**	**0.6647**	**0.5016**	**0.3914**
Taobao	DeepFM	0.5303	0.1305	0.2999	0.4261	0.1983	0.2519	0.1089	0.0983	0.1010	0.2497	0.3736	0.1097	0.1461	0.1508
	ESMM	0.5043	0.1088	0.2838	0.4152	0.1823	0.2376	0.1023	0.0952	0.0847	0.2366	0.3640	0.0918	0.1377	0.1464
	MMOE	0.5044	0.1068	0.2777	0.4097	0.1781	0.2333	0.1003	0.0941	0.0824	0.2317	0.3591	0.0896	0.1349	0.1445
	BERT4Rec	0.5039	0.1020	0.2767	0.4168	0.1758	0.2350	0.1004	0.0964	0.0783	0.2303	0.3670	0.0853	0.1347	0.1480
	ESM2	0.5282	0.1289	0.3036	0.4304	0.1999	0.2543	0.1103	0.0994	0.0993	0.2529	0.3779	0.1080	0.1479	0.1525
	ESCM^2-DR	0.5127	0.1109	0.2853	0.4192	0.1831	0.2401	0.1030	0.0965	0.0850	0.2371	0.3677	0.0925	0.1384	0.1482
	ESCM^2-IPS	0.5128	0.1148	0.2893	0.4218	0.1876	0.2436	0.1050	0.0971	0.0893	0.2413	0.3700	0.0968	0.1409	0.1492
	DCMT	0.5093	0.1102	0.2857	0.4122	0.1833	0.2368	0.1033	0.0947	0.0847	0.2370	0.3605	0.0922	0.1386	0.1454
	CL4CVR	0.5106	0.1112	0.2884	0.4197	0.1852	0.2411	0.1048	0.0970	0.0851	0.2403	0.3684	0.0927	0.1405	0.1489
	DDPO	<u>0.5626</u>	<u>0.1378</u>	<u>0.3262</u>	<u>0.4536</u>	<u>0.2154</u>	<u>0.2705</u>	<u>0.1186</u>	<u>0.1051</u>	<u>0.1073</u>	<u>0.2739</u>	<u>0.4007</u>	<u>0.1162</u>	<u>0.1595</u>	<u>0.1614</u>
	NISE	0.5068	0.1096	0.2889	0.4182	0.1848	0.2400	0.1050	0.0969	0.0841	0.2405	0.3679	0.0916	0.1408	0.1486
	EVI	0.5061	0.1081	0.2784	0.4095	0.1789	0.2339	0.1006	0.0942	0.0837	0.2320	0.3583	0.0909	0.1353	0.1446
	PARS (Ours)	**0.8953**	**0.1768**	**0.3476**	**0.4572**	**0.2444**	**0.2926**	**0.1282**	**0.1070**	**0.1407**	**0.2961**	**0.4069**	**0.1514**	**0.1726**	**0.1642**

Table 2: Performance comparison across three datasets. Best results are in **bold**, second best are <u>underlined</u>.

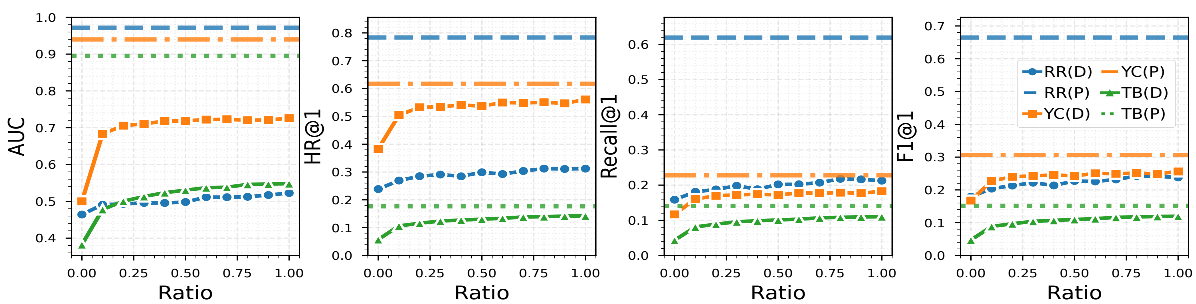

Figure 3: Performance comparison across multiple evaluation metrics with varying the ratios of using positive acquisition labels on three e-commerce datasets. (Abbreviations: RR=RetailRocket, YC=YooChoose, TB=Taobao, D=DeepFM, P=PARS.)

- **DDPO** (Su et al. 2024): Doubly-robust estimator with Dual Propensity Optimization that combines inverse propensity weighting with direct prediction.
- **NISE** (Huang et al. 2024a): Negative sampling with Importance Sampling Estimator that addresses exposure bias through importance-weighted learning.
- **EVI** (Fei, Zhang, and Li 2025): An entire-space variational information exploitation framework that generates unbiased pseudo labels via a conditional CVR teacher and transfers non-click information using variational inference and logit distillation to improve CVR prediction.

These baseline methods represent the current state-of-the-art in CVR prediction, covering diverse approaches that include multi-task learning, causal inference, deep feature in-

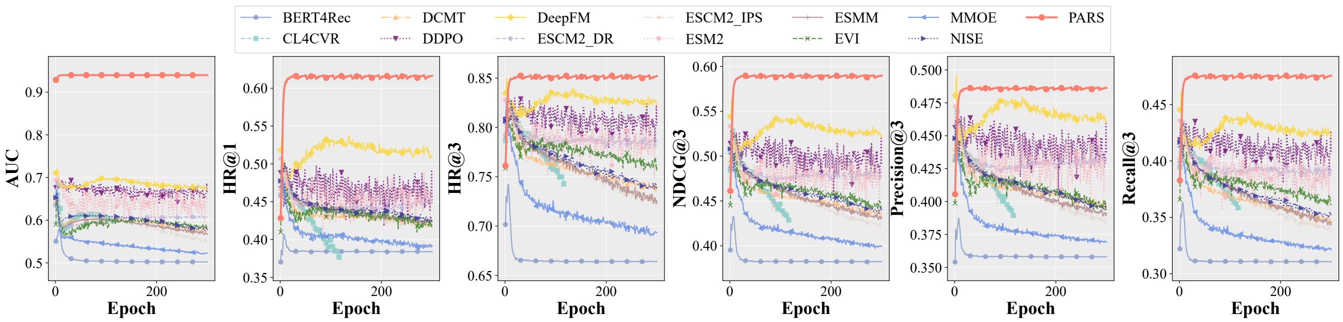

Figure 4: Comparison of our method and baselines on the YooChoose dataset across six metrics (AUC, HR@1, HR@3, NDCG@3, Precision@3, Recall@3) over training epochs.

teraction modeling, and various debiasing techniques. Compared against this comprehensive set of baselines, we can thoroughly evaluate the effectiveness of PARS. Our implementation is publicly available at the following link[4].

Impact of Label Availability on Baseline Performance We investigate how label availability affects baseline model performance. Ablation studies are conducted by providing baseline training models with varying ratios of ground-truth labels. Note that when the ratio equals 50%, it means that the baseline methods use 50% of the positive acquisition labels and set the rest positive acquisition labels to zeros. Figure 3 illustrates the performances of DeepFM using different label ratios on all three datasets. Remarkably, PARS without using item acquisition labels outperforms DeepFM using 100% item acquisition labels. More experiments are postponed to supplementary materials.

The results reveal a critical finding: baseline performance deteriorates significantly when label availability drops below 10%, with AUC scores declining by up to 35% on YooChoose and 42% on Taobao. This dramatic performance drop underscores the heavy reliance of traditional CVR models on labeled data.

Overall Comparative Analysis Although PARS outperforms most CVR baselines using 100% item acquisition labels, its performance on Taobao is not significant because Taobao has a very high number of positive item acquisition labels, i.e., 75.45% of the user interaction sequences contain at least one item marked as purchased. Table 2 shows that, under different metrics, PARS without using item acquisition labels remarkably outperforms all baselines using 50% positive item acquisition labels, showing the effectiveness of PARS. In the supplementary materials, we show that PARS without using any item acquisition labels outperforms most of the CVR baselines using 100% item acquisition labels.

Training Stability and Convergence Analysis To evaluate the training dynamics and convergence behavior of our proposed methods, we conduct extensive experiments tracking model performance throughout the training process. Figure 4 presents a comprehensive comparison between our approach and baseline methods on the YooChoose datasets (comparisons on RetailRocket and Taobao datasets are put in the supplementary materials due to space limit), evaluating six key metrics: AUC, HR@1, HR@3, NDCG@3, Precision@3, and Recall@3 over 300 training epochs. The results reveal several important insights about our label-free learning framework. First, PARS demonstrates remarkable training stability across all metrics and datasets. While baseline methods exhibit considerable fluctuations during training, particularly in the initial 50 epochs, PARS maintains a smooth and consistent improvement trajectory. This stability is particularly pronounced in ranking metrics such as NDCG@3 and HR@3, where PARS shows monotonic improvement without the oscillations commonly observed in baselines. Second, our approach achieves faster convergence compared to most baselines. On the YooChoose dataset, PARS reaches 95% of its final performance within 50 epochs, while baselines such as DDPO and DeepFM require approximately 150-200 epochs to stabilize. This efficient convergence is attributed to our principled learning framework that leverages inherent data structures rather than relying on potentially noisy label signals. Furthermore, the consistent performance across all six evaluation metrics demonstrates the robustness of PARS. Unlike some baselines that show trade-offs between different metrics (e.g., improving AUC while degrading Precision@3), PARS maintains balanced performance improvements across both ranking and classification objectives. This comprehensive superiority validates the effectiveness of PARS.

Conclusion

The paper presents a novel and promising RS setting, where only user browsing histories are available without any item acquisition labels. An effective algorithm termed as partial acquisition recommender system (PARS) has been proposed. The key idea is to use maksed language modeling and partial label learning to help extract reliable user and item embeddings to predict item acquisition probabilities. Experiments show that PARS without using item acquisition labels outperforms baselines using item acquisition labels. In the future, we will extend the work in the following aspects. First, build a theoretical foundation of the proposed method. Second, how to effectively incorporate a small amount of item acquisition labels when available.

[4] https://github.com/bjtu-lucas-nlp/PARS

Acknowledgments

This work is supported by ARC Discovery Project DP220102635 "Robust Meta Learning for Risk-aware Recommender Systems."

References

Claypool, M.; Brown, D.; Le, P.; and Waseda, M. 2001. Inferring user interest. *IEEE Internet Computing*, 5(6): 32–39.

Cour, T.; Sapp, B.; and Taskar, B. 2011. Learning from partial labels. *The Journal of Machine Learning Research*, 12: 1501–1536.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.

Dwork, C. 2006. Differential privacy. In *International colloquium on automata, languages, and programming*, 1–12. Springer.

Fei, K.; Zhang, X.; and Li, J. 2025. Entire-Space Variational Information Exploitation for Post-Click Conversion Rate Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 11654–11662.

Feng, L.; Lv, J.; Han, B.; Xu, M.; Niu, G.; Geng, X.; An, B.; and Sugiyama, M. 2020. Provably consistent partial-label learning. *Advances in neural information processing systems*, 33: 10948–10960.

Gui, J.; Chen, T.; Zhang, J.; Cao, Q.; Sun, Z.; Luo, H.; and Tao, D. 2024. A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9052–9071.

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*.

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, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; and Yang, Q. 2024b. Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9387–9406.

Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, 197–206. IEEE.

Liang, F.; Pan, W.; and Ming, Z. 2021. Fedrec++: Lossless federated recommendation with explicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4224–4231.

Liu, Z.; Chen, Y.; Li, J.; Yu, P. S.; McAuley, J.; and Xiong, C. 2021. Contrastive self-supervised sequential recommendation with robust augmentation. *arXiv preprint arXiv:2108.06479*.

Lu, J.; Wu, D.; Mao, M.; Wang, W.; and Zhang, G. 2015. Recommender system application developments: a survey. *Decision support systems*, 74: 12–32.

Lu, K.; Lu, J.; Xu, H.; Guo, K.; Zhang, Q.; Lin, H.; Grosser, M.; Zhang, Y.; and Zhang, G. 2025. Genomics-Enhanced Cancer Risk Prediction for Personalized LLMs-Driven Healthcare Recommender Systems. *ACM Transactions on Information Systems*.

Lu, K.; Zhang, Q.; Hughes, D.; Zhang, G.; and Lu, J. 2024. Amt-cdr: A deep adversarial multi-channel transfer network for cross-domain recommendation. *ACM Transactions on Intelligent Systems and Technology*, 15(4): 1–26.

Lv, J.; Xu, M.; Feng, L.; Niu, G.; Geng, X.; and Sugiyama, M. 2020. Progressive identification of true labels for partial-label learning. In *international conference on machine learning*, 6500–6510. PMLR.

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.

Núñez-Valdez, E. R.; Quintana, D.; Crespo, R. G.; Isasi, P.; and Herrera-Viedma, E. 2018. A recommender system based on implicit feedback for selective dissemination of ebooks. *Information Sciences*, 467: 87–98.

Ouyang, W.; Dong, R.; Zhang, X.; Guo, C.; Luo, J.; Liu, X.; and Du, Y. 2023. Contrastive Learning for Conversion Rate Prediction. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1909–1913.

Rajendran, D. P. D.; and Sundarraj, R. P. 2021. Using topic models with browsing history in hybrid collaborative filtering recommender system: Experiments with user ratings. *International Journal of Information Management Data Insights*, 1(2): 100027.

Resnick, P.; and Varian, H. R. 1997. Recommender systems. *Communications of the ACM*, 40(3): 56–58.

Roy, D.; and Dutta, M. 2022. A systematic review and research perspective on recommender systems. *Journal of Big Data*, 9(1): 59.

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.

Sun, F.; Liu, J.; Wu, J.; Pei, C.; Lin, X.; Ou, W.; and Jiang, P. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*, 1441–1450.

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, S.; Hu, L.; Wang, Y.; Cao, L.; Sheng, Q.; and Orgun, M. 2019. Sequential recommender systems: Challenges, progress and prospects. In *Twenty-Eighth International Joint Conference on Artificial Intelligence {IJCAI-19}*. International Joint Conferences on Artificial Intelligence Organization.

Wen, H.; Zhang, J.; Lin, Q.; Yang, K.; and Huang, P. 2019. Multi-level deep cascade trees for conversion rate prediction in recommendation system. In *proceedings of the AAAI conference on artificial intelligence*, volume 33, 338–345.

Wen, H.; Zhang, J.; Wang, Y.; Lv, F.; Bao, W.; Lin, Q.; and Yang, K. 2020. Entire space multi-task modeling via post-click behavior decomposition for conversion rate prediction. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 2377–2386.

Xie, X.; Sun, F.; Liu, Z.; Wu, S.; Gao, J.; Zhang, J.; Ding, B.; and Cui, B. 2022. Contrastive learning for sequential recommendation. In *2022 IEEE 38th international conference on data engineering (ICDE)*, 1259–1273. IEEE.

Xu, M.; Lian, Z.; Feng, L.; Liu, B.; and Tao, J. 2023. ALIM: adjusting label importance mechanism for noisy partial label learning. *Advances in Neural Information Processing Systems*, 36: 38668–38684.

Xu, N.; Lv, J.; and Geng, X. 2019. Partial label learning via label enhancement. In *Proceedings of the AAAI Conference on artificial intelligence*, volume 33, 5557–5564.

Xu, N.; Qiao, C.; Geng, X.; and Zhang, M.-L. 2021. Instance-dependent partial label learning. *Advances in Neural Information Processing Systems*, 34: 27119–27130.

Ye, S.; and Lu, J. 2023. Sequence unlearning for sequential recommender systems. In *Australasian Joint Conference on Artificial Intelligence*, 403–415. Springer.

Ye, S.; and Lu, J. 2024. Robust recommender systems with rating flip noise. *ACM Transactions on Intelligent Systems and Technology*, 16(1): 1–19.

Ye, S.; Lu, J.; and Zhang, G. 2025. Towards safe machine unlearning: A paradigm that mitigates performance degradation. In *Proceedings of the ACM on Web Conference 2025*, 4635–4652.

Zhang, M.-L.; Zhou, B.-B.; and Liu, X.-Y. 2016. Partial label learning via feature-aware disambiguation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1335–1344.

Zhou, Z.-H. 2018. A brief introduction to weakly supervised learning. *National science review*, 5(1): 44–53.

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. IEEE.