The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Counterfactual Task-augmented Meta-learning for Cold-start Sequential Recommendation

Zhiqiang Wang[1], **Jiayi Pan**[1], **Xingwang Zhao**[1*], **Jianqing Liang**[1], **Chenjiao Feng**[2], **Kaixuan Yao**[1]
[1]Shanxi Taihang Laboratory, School of Computer and Information Technology, Shanxi University, Taiyuan, China
[2]College of Applied Mathematics, Shanxi University of Finance and Economics, Taiyuan, China
wangzq@sxu.edu.cn, panjiayi0618@foxmail.com, {zhaoxw, liangjq}@sxu.edu.cn,
fengcj@sxufe.edu.cn, ykx@sxu.edu.cn

Abstract

Cold-start sequential recommendation, where user interaction histories are sparse or minimal, remains a significant challenge in recommendation systems. Current meta-learning-based approaches rely heavily on the interaction histories of regular users to construct meta-tasks, aiming to acquire prior knowledge for cold-start adaptation. However, these methods often fail to account for preference discrepancies between regular and cold-start users, leading to biased preference modeling and suboptimal recommendations. To address this issue, we propose a novel counterfactual task-augmented meta-learning method for cold-start sequential recommendations. Our approach intervenes in user interaction histories to create counterfactual sequences that simulate potential but unrealized user behaviors, establishing counterfactual tasks within a meta-learning framework. Additionally, we aggregate meta-path neighbors to uncover latent relationships between items, enabling more detailed and accurate modeling of user preferences. Moreover, by integrating real and counterfactual task losses, we jointly optimize the model through a combination of global and local updates, enhancing its adaptability to cold-start scenarios. Extensive experiments demonstrate that our method significantly outperforms existing state-of-the-art techniques, achieving superior results in cold-start sequential recommendation tasks.

Introduction

Sequential recommendation (Sun et al. 2019; Kang and McAuley 2018) is a critical task in recommender systems, utilizing user behavior sequences such as browsing, purchasing, and clicking to infer preferences and predict future interactions. The challenge lies in accurately modeling both short-term and long-term shifts in user interests. In cold-start scenarios, where historical interaction data is sparse, the problem of providing effective recommendations for new users or items becomes particularly acute. Despite the introduction of numerous methods for cold-start sequential recommendation, significant challenges persist, particularly in dynamically capturing user preferences with limited interaction data.

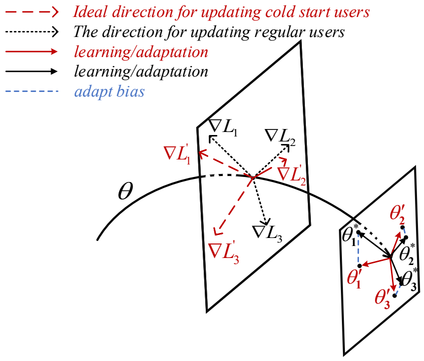

Figure 1: Optimization divergence between regular and cold-start user updates in a meta-learning framework.

Existing approaches to cold-start recommendation can be broadly classified into two categories: methods utilizing auxiliary information and those leveraging limited interaction data (Du et al. 2022). The former enhances recommendation performance by incorporating item attributes, knowledge graphs, or multi-behavioral data. For instance, DropoutNet (Volkovs, Yu, and Poutanen 2017) reduces dependency on ID embeddings by utilizing auxiliary data to refine user and item representations, whereas KGNN-LS (Wang et al. 2019) leverages structured knowledge from graphs to address cold-start challenges. However, these methods based on auxiliary information often encounter challenges, not only due to data privacy concerns but also because such information is not always readily available or sufficiently comprehensive, particularly for cold-start users. The latter category focuses on leveraging limited interaction data, with meta-learning methods being employed to rapidly adapt models to cold-start scenarios. For example, metaCSR (Huang et al. 2022) initializes model parameters based on regular user behavior, facilitating quick adaptation with minimal data, while Meta TL (Wang, Ding, and Caverlee 2021) learns a universal initialization parameter via meta-learning algorithms, fine-tuning the model for each cold-start user. Nevertheless, these methods tend to over-rely on the interaction data of regular users, thereby neglecting the unique behavioral patterns of cold-start users. This dependency can introduce biases into the models, limiting their ability to capture unobserved potential sequences and resulting in suboptimal recommendations.

---

[*] Corresponding author.

Figure 1 illustrates the optimization divergence between updates based on regular users and the ideal updates for cold-start users within a meta-learning framework. The red dashed line indicates the ideal update direction for cold-start users, while the black dashed line represents the update direction derived from regular users. The deviation in direction and distance when relying on regular user data leads to suboptimal adjustments for cold-start users. Traditional methods predominantly use regular user data (black dashed arrows) to adapt model parameters (θ) through observed interactions (solid black arrows), which often results in biased updates.

In cold-start scenarios, the inherent scarcity of interaction data necessitates the optimal utilization of all available information. While observed user interactions provide direct insights into user preferences, unobserved potential sequences—referred to as counterfactual sequences—represent hypothetical interactions that could have occurred under different conditions. Capturing and leveraging these counterfactual sequences effectively expands the hypothesis space for preference modeling, offering a viable approach to mitigate biases associated with reliance solely on regular user data.

To address these challenges, we propose a novel counterfactual task-augmented meta-learning (CTM) method for cold-start sequential recommendation. Our approach intervenes in user interaction histories to generate counterfactual sequences, constructing counterfactual tasks within a meta-learning framework to more accurately model user preferences. Additionally, we introduce a meta-path neighbor aggregation strategy that uncovers latent relationships between items, thereby enhancing the granularity of preference modeling. By jointly optimizing both real and counterfactual task losses, our model significantly improves adaptability to cold-start scenarios, resulting in enhanced recommendation accuracy. The primary contributions of this work are summarized as follows:

- We propose a counterfactual task-augmented meta-learning method for cold-start sequential recommendation, leveraging meta-learning to capture user interactions and address preference discrepancies between regular and cold-start users.
- We design a meta-path neighbor aggregation mechanism to model intricate inter-item relationships, enhancing the representation of latent user preferences.
- We validate the effectiveness of the CTM method through extensive experiments on three public datasets, demonstrating superior performance over state-of-the-art methods in cold-start scenarios.

Related Work

The cold-start recommendation problem has garnered significant attention in recommender systems, leading to the development of a wide range of approaches (Wang et al. 2023; Wu et al. 2022). These approaches can be broadly categorized into those leveraging auxiliary information and those relying on limited interaction data, spanning content-based, collaborative filtering (Xu et al. 2023), deep learning (Wu et al. 2023; Jing et al. 2024), and meta-learning (Du et al. 2023; Shu, Chung, and Lin 2023) paradigms.

Auxiliary information-based methods enhance recommendation models by incorporating additional data, such as user demographics and item attributes. For example, DropoutNet (Volkovs, Yu, and Poutanen 2017) reduces dependency on ID embeddings by leveraging content features, while AGNN (Qian et al. 2023) employs a variational autoencoder to aggregate multimodal attributes in cold-start scenarios. Knowledge graph-based approaches like KGNN-LS (Wang et al. 2019) enrich user and item representations by integrating graph neural networks with structured knowledge, enabling richer feature aggregation. Cross-domain methods, such as EMCDR (Man et al. 2017), transfer knowledge from auxiliary domains to improve target domain recommendations, with models like IHGNN (Cai et al. 2023) retaining relational information through heterogeneous graphs. Multi-behavior-based methods, exemplified by MBGCN (Jin et al. 2020), incorporate various user behaviors to uncover deeper insights into preferences. Despite their effectiveness, these methods often face challenges related to data scarcity, inadequacy, or privacy issues, which limit the availability of necessary information and potentially reduce recommendation efficacy.

In contrast, meta-learning approaches have emerged as powerful tools for cold-start scenarios where interaction data is sparse, enabling rapid adaptation by capturing task commonalities. Metric-based meta-learning methods, such as MWUF (Zhu et al. 2021b), enhance cold-start embeddings by refining feature spaces, while model-based approaches like Mecos (Zheng et al. 2021) dynamically generate model parameters for swift adaptation. Optimization-based methods, such as MAML (Finn, Abbeel, and Levine 2017), focus on extracting transferable representations across tasks, with models like MeLU (Lee et al. 2019) and metaCSR (Huang et al. 2022) employing few-shot learning to adapt quickly to new users or items. TMCDR (Zhu et al. 2021a) integrates transfer learning with meta-learning for cross-domain recommendation. However, relying on regular user data to simulate cold-start scenarios can introduce biases, compromising accuracy for true cold-start users. This work focuses on optimization-based meta-learning approaches designed to effectively address the challenges associated with cold-start recommendations in short sequence scenarios.

Research on counterfactual data augmentation holds promise in recommendation systems, particularly for addressing cold-start issues and improving system robustness and fairness. Wang et al. proposed the CASR (Wang et al. 2021), which utilizes counterfactual theory to generate new user behavior sequence data by employing heuristic sampling, data-based sampling, and model-based sampling to create data beneficial for model training. Peng et al. introduced the CALRec (Peng, Song, and Liu 2023), which combines the strengths of active learning and generative data augmentation. By pretraining a sampling model and an anchor model, the framework extracts counterfactual data from the sampling model and applies an active learning strategy to filter the best counterfactual data to enhance model performance. Compared to traditional random data augmenta-

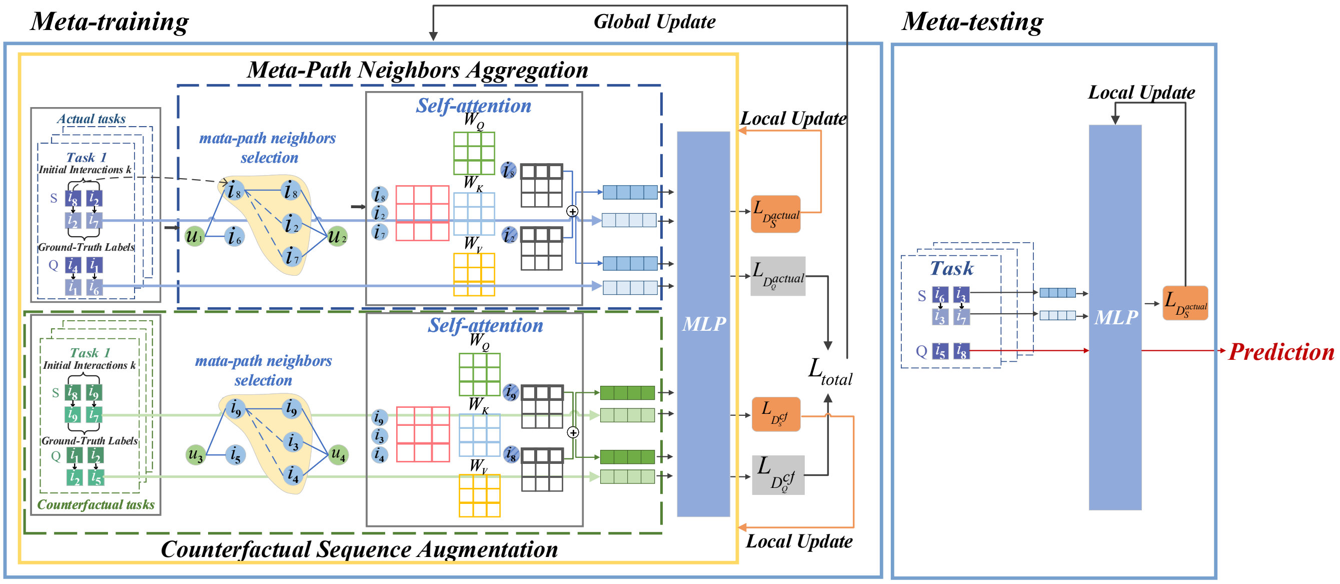

Figure 2: CTM Framework.

tion methods, counterfactual data augmentation offers more meaningful data variations that the model can better understand and utilize, thereby improving recommendation system performance and interpretability.

Method

CTM Framework

The proposed CTM method is designed to mitigate biases in cold-start user preference modeling by leveraging recommendation patterns learned from regular user interactions. As shown in Figure 2, the CTM framework includes three key components: counterfactual sequence augmentation, meta-path neighbor aggregation, and meta-learning optimization.

The counterfactual sequence augmentation component generates counterfactual sequences by intervening in users' historical interaction sequences, expanding the hypothesis space for preference modeling. The meta-path neighbor aggregation component enhances the representation of latent user preferences by aggregating meta-path neighbors related to user interactions. Finally, the meta-learning optimization component jointly optimizes both real and counterfactual task losses through a combination of global and local updates, ensuring robust adaptation in cold-start scenarios. These components are detailed in following subsections.

Counterfactual Sequence Augmentation

The core objective of the counterfactual sequence augmentation module is to simulate alternative user behaviors by intervening in the historical interaction sequence R_{actual}. Specifically, this module introduces controlled modifications to key items within the sequence, preserving the overall structural integrity while introducing sufficient variability to generate a counterfactual sequence R_{cf}. It aims to capture potential shifts in user preferences that could result from different item interactions.

The intervention function $g(R_{actual}, \sigma)$ is to modify a user's historical interaction sequence. Given the sequence $R_{actual} = \{i_1^u, i_2^u, \ldots, i_k^u, \ldots, i_t^u\}$, the sequence after intervention is:

$$\begin{aligned} R_{cf} &= g(R_{actual}, \sigma) \\ &= \{i_1^u, (i_2)^1, \ldots, (i_k)^e, \ldots, (i_{t-1})^{t\cdot\sigma}, i_t^u\} \end{aligned} \tag{1}$$

where σ represents the degree of intervention on the original sequence R_{actual}, defined as the proportion of items in the sequence randomly replaced. The number of replaced items is $t \cdot \sigma$. In this sequence, $(i_2)^1$ and $(i_k)^e$ represent items in R_{actual} selected to replace i_2^u and i_k^u, respectively, with e indicating the position of the e-th replaced item. By generating diverse counterfactual sequences through interventions in user interaction sequences, the model expands its hypothesis space for cold-start preferences beyond the constraints of regular supervision.

Within the meta-learning framework, the implementation begins by sampling a task T_n for each user u_j, where $T_n = \{R_{\text{actual}} \cup R_{\text{cf}}\}$, consisting of pairs of actual and counterfactual sequences. The subsequent item in each sequence is treated as the target label, with attention scores for each item computed using a self-attention mechanism. The representation of an item i_n is then computed as:

$$i_n = \sum_{m=1}^{T} \text{Attention}(i_m, i_n) \cdot i_n, \tag{2}$$

where i_m represents a neighboring item of i_n, and $\text{Attention}(i_m, i_n)$ denotes the computed attention score reflecting the relevance of i_m to i_n. Detailed calculations of these attention scores are elaborated in Section 3.3, "Meta-path Neighbor Aggregation".

Subsequently, each actual sequence $E_u^{\text{actual}} = \{i_1, i_2, \ldots, i_k\} \in \mathbb{R}^{T \times d}$ is projected onto a vector $S_u^{\text{actual}} \in \mathbb{R}^d$, and each counterfactual sequence $E_u^{\text{cf}} = \{i_1, i_2, \ldots, i_k\} \in \mathbb{R}^{T \times d}$ is similarly projected onto a vector $S_u^{\text{cf}} \in \mathbb{R}^d$ through fully connected layers:

$$S_u^{\text{actual}} = \text{ReLU}(W_1 E_u^{\text{actual}} + b_1), \tag{3}$$

$$S_u^{\text{cf}} = \text{ReLU}(W_2 E_u^{\text{cf}} + b_2), \tag{4}$$

where S_u^{actual} and S_u^{cf} denote the embedded representations of the actual and counterfactual sequences, respectively. The interaction probability between user u and item i is predicted using a multilayer perceptron (MLP) with three layers, as follows:

$$p_{u,i} = \sigma(W_2 \,\text{ReLU}(W_1 \,\text{ReLU}(W_0[S_u, i] + b_0) + b_1) + b_2), \tag{5}$$

where $[S_u, i]$ is the concatenated input of the sequence representation S_u (either S_u^{actual} or S_u^{cf}) and the item i. Here, W_0, W_1, and W_2 are the weight matrices, and b_0, b_1, and b_2 are the bias terms for each corresponding layer. The function $\sigma(\cdot)$ represents the activation function applied to the output.

Meta-path Neighbors Aggregation

This module aims to model global user preferences by uncovering deeper relationships between users and items. A meta-path is a path in a heterogeneous graph comprising various types of nodes and edges. In a "user–item" bipartite graph, meta-paths represent user-item relationships. In a heterogeneous user-item interaction graph, capturing meaningful inter-item relationships is crucial, especially in cold-start scenarios where interaction data is sparse. To mitigate this sparsity, we leverage regular user interaction data to enhance preference modeling for cold-start users.

In a user-item interaction graph, items connected by the same user indicate similar preferences. Specifically, if two users have interacted with the same item, they likely share preferences for other items as well. For example, the path $(u_1 \rightarrow i_8 \rightarrow u_2)$ in Figure 3 indicates that both users u_1 and u_2 interacted with item i_8, implying that items i_2, i_7, and i_4, as meta-path neighbors of i_8, may reveal hidden preferences.

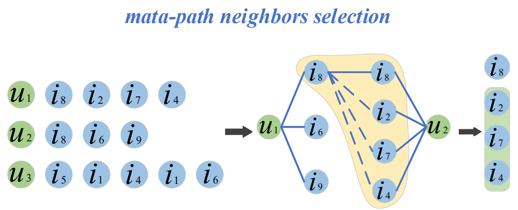

Figure 3: Example of meta-path neighbor selection

In this approach, each item ID is initially mapped to a learnable embedding matrix to obtain item embeddings. A self-attention mechanism is then applied as the aggregation function:

$$i = \text{SelfAttention}([t_i, t_1^{nb}, \ldots, t_s^{nb}]), \tag{6}$$

where t_i represents the embedding of item i, and s denotes the number of neighbors for item i. The self-attention mechanism allows each item to account for the influence of its neighbors, capturing long-range dependencies:

$$Q = IW^Q,\ K = IW^K,\ V = IW^V, \tag{7}$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right) V, \tag{8}$$

where W^Q, W^K, and W^V are the weight matrices, and d is the dimensionality of the embeddings. This method enables fine-grained modeling of complex relationships, enhancing the exploration of user preferences.

Within the meta-learning model proposed in this paper, the meta-path neighbor aggregation module plays a critical role. By incorporating global structural information from regular users, this module enriches item representations, facilitating a deeper understanding of item relationships. This enhancement enables subsequent modules to improve the accuracy of the recommendation for cold-start users. Integration with the meta-learning framework ensures effective generalization across tasks, thereby enhancing performance in cold-start scenarios.

Meta-learning Optimization

The proposed CTM model optimizes real and counterfactual task losses using global and local updates within a meta-learning framework. During meta-task training, let $p(T)$ denote the distribution over all tasks. Each task T_i involves N users, where each user selects s_1 behavior sequences for the support set D_S^{actual} and s_2 sequences for the query set D_Q^{actual}. The Counterfactual Sequence Augmentation Module generates corresponding counterfactual sets D_S^{cf} and D_Q^{cf}, unified as:

$$D_S = D_S^{actual} \cup D_S^{cf}, \quad D_Q = D_Q^{actual} \cup D_Q^{cf}. \tag{9}$$

Item representations are learned through the Meta-Path Neighbor Aggregation Module, resulting in user preference embeddings. Model parameters are denoted as $f_{(\phi,\Theta)}$, where Θ pertains to the task-level parameters, and ϕ to the global parameters. These parameters are initialized randomly.

The Counterfactual Training Module predicts the probabilities $p_{(u,i)}^{actual}$ and $p_{(u,i)}^{cf}$, with loss functions defined as:

$$\mathcal{L}_{D_Q^{actual}} = \sum_u \sum_i \left[-\log \left(p_{(u,i)}^{actual} - p_{(u,i^-)}^{actual} \right) \right], \tag{10}$$

$$\mathcal{L}_{D_Q^{cf}} = \sum_u \sum_i \left[-\log \left(p_{(u,i)}^{cf} - p_{(u,i^-)}^{cf} \right) \right], \tag{11}$$

$$\mathcal{L}_{total} = \lambda \mathcal{L}_{D_Q^{actual}} + (1 - \lambda) \mathcal{L}_{D_Q^{cf}}, \tag{12}$$

where $\mathcal{L}_{D_Q^{actual}}$ and $\mathcal{L}_{D_Q^{cf}}$ represent the losses for the actual and counterfactual sequences, respectively, and L_{total} is the weighted sum of these losses.

Algorithm 1: Meta-Training of CTM

Input: Distribution over tasks $p(T)$, step size hyperparameters α, β, parameters ϕ, Θ

Output: Optimized parameter Θ^*

```
Randomly initialize $\phi, \Theta$
while not converged do
    Sample a batch of tasks $T_i \sim p(T)$
    for all tasks $T_i^{actual}, T_i^{cf}$ do
        Sample $s_1$ sequences as $D_S^{actual}$ from $T_i^{actual}$
        Sample $s_2$ sequences as $D_S^{cf}$ from $T_i^{cf}$
        Evaluate $\nabla_\Theta \mathcal{L}(f_{(\phi,\Theta)})$ using $D_S^{actual}$ and $D_S^{cf}$
        Local update:
        Compute adapted parameters via gradient descent:
        $\Theta^* \leftarrow \Theta - \alpha \nabla_\theta \mathcal{L}(f_{(\phi,\Theta)})$
    end for
    Global update:
    Sample $s_3$ sequences as $D_Q^{actual}$ from $T_i^{actual}$
    Sample $s_4$ sequences as $D_Q^{cf}$ from $T_i^{cf}$
    Evaluate $\nabla_\Theta \mathcal{L}_{total}(f_{(\phi,\Theta^*)})$ using $D_Q^{actual}$ and $D_Q^{cf}$
    Update $\phi$ :
    $\phi \leftarrow \phi - \beta \nabla_\phi \sum_{T_n \sim p(T)} \mathcal{L}_{total}(f_{(\phi,\Theta^*)})$
end while
```

During the inner-loop optimization phase, the support set D_S is used to update Θ via stochastic gradient descent, aiming to minimize the total loss:

$$\Theta^* = \Theta - \alpha \nabla_\Theta \mathcal{L}(f_{(\phi,\Theta)}), \tag{13}$$

where α is the learning rate. This step refines the initialization of item representations to obtain personalized user preference embeddings.

In the outer-loop optimization phase, the model $f_{(\phi,\Theta^*)}$ is updated using the query set D_Q derived from the training tasks, with the objective of minimizing the total loss:

$$\min_{\phi,\Theta} \sum_{T_n \sim p(T)} \mathcal{L}(f_{(\phi,\Theta^*)}). \tag{14}$$

The final updated parameters after stochastic gradient descent optimization are:

$$\phi \leftarrow \phi - \beta \nabla_\phi \sum_{T_n \sim p(T)} \mathcal{L}_{total}(f_{(\phi,\Theta^*)}). \tag{15}$$

This iterative process updates the model parameters through meta-learning, allowing for rapid adaptation to cold-start users and enhancing recommendation accuracy. The complete algorithm for the CTM model is detailed in Algorithm 1.

Experiments

Datasets and Experimental Setup

Datasets. To evaluate the effectiveness of model, experiments were run on three benchmark datasets, chosen for their diversity and relevance. Dataset statistics are in Table 1.

Dataset	Users	Items	Interactions
Electronics	22,685	20,712	10,000
MovieLens 100k	1,000	1,700	100,000
KuaiRec	1,411	3,327	4,676,570

Table 1: Summary of Datasets

- **Electronics.** This dataset is a subset of the Amazon Review Data (He and McAuley 2016), widely utilized in recommendation system research. It comprises rating records from 22,685 users on 20,712 electronics products, covering interactions from May 1996 to July 2014.
- **MovieLens 100k.** A well-known dataset in the recommendation systems domain (Harper and Konstan 2015), curated by the GroupLens research group. It contains 100,000 ratings from 943 users on 1,682 movies, including timestamps, making it a valuable benchmark for algorithm evaluation.
- **KuaiRec.** A publicly available dataset for recommendation research (Gao et al. 2022), provided by the University of Science and Technology of China and Kuaishou Technology Co., Ltd. It includes 4,676,570 interaction records from 1,411 users on 3,327 videos, enriched with user demographics, social relationships, and video tags.

Evaluation Metrics. To evaluate the proposed method, each cold-start user is linked to a benchmark item, and standard metrics are used to assess performance:

- **MRR.** MRR evaluates the ranking effectiveness by considering the position of the first relevant item in the ranked list. It is computed as:

  $$\text{MRR} = \frac{1}{Q} \sum_{q=1}^{Q} \frac{1}{\text{rank}_{\text{relevant}}(q)}, \tag{16}$$

  where Q is the total number of queries, and $\text{rank}_{\text{relevant}}(q)$ is the rank of the first relevant item for query q.
- **Hits@k.** Hits@k measures the frequency at which a relevant item appears within the top k positions in the ranked list. It is defined as:

  $$\text{Hits@k} = \frac{1}{N} \sum_{i=1}^{N} I(\text{rank}_i \leq k), \tag{17}$$

  where N is the user count, rank_i is the relevant item's rank for user i, and $I(\cdot)$ is an indicator function: 1 if true, 0 if false.
- **NDCG@k.** NDCG@k measures ranking quality by comparing the top k items' cumulative gain to an ideal ranking. It is calculated as:

  $$\text{NDCG@k} = \frac{\text{DCG@k}}{\text{IDCG@k}}, \tag{18}$$

  where the Discounted Cumulative Gain (DCG@k) is:

  $$\text{DCG@k} = \sum_{i=1}^{k} \frac{2^{\text{rel}_i} - 1}{\log_2(i+1)}, \tag{19}$$

  and IDCG@k is the maximum possible DCG@k. rel_i denotes the relevance score of the item at rank i.

Dataset	Electronics				MovieLens 100k				KuaiRec			
Method	MRR	NDCG@5	Hits@5	Hits@1	MRR	NDCG@5	Hits@5	Hits@1	MRR	NDCG@5	Hits@5	Hits@1
Wide&Deep	0.181	0.176	0.273	0.082	0.167	0.197	0.226	0.105	0.179	0.188	0.215	0.093
TransRec	0.294	0.208	0.428	0.183	0.234	0.215	0.366	0.144	0.223	0.179	0.260	0.115
SASRec	0.318	0.323	0.409	0.193	0.284	0.247	0.378	0.142	0.196	0.174	0.263	0.117
MeLU	0.265	0.309	0.412	0.172	0.241	0.298	0.405	0.124	0.334	0.402	0.483	0.223
Meta TL	0.348	0.317	0.480	0.217	0.293	0.317	0.458	0.145	0.371	0.393	0.603	0.191
CTM	**0.363**	**0.365**	**0.498**	**0.238**	**0.347**	**0.363**	**0.539**	**0.224**	**0.520**	**0.585**	**0.846**	**0.305**

Table 2: Performance Comparison on Different Datasets

Comparison Methods. To evaluate the proposed model's performance, we compare it against several state-of-the-art baselines commonly used in recommendation systems:

- **Wide&Deep** (Cheng et al. 2016). This model combines linear models (Wide) and deep neural networks (Deep) to simultaneously capture both direct feature associations and complex feature interactions, offering superior performance over models using either component alone.
- **TransRec** (He, Kang, and McAuley 2017). TransRec models user preferences as translation vectors in an embedding space, enabling effective sequence-based recommendations by capturing the relationships between items and user transitions.
- **SASRec** (Kang and McAuley 2018). A sequence-based recommendation model that employs self-attention to capture long-term dependencies in user behavior sequences, identifying the most relevant next item by focusing on key elements in the sequence.
- **MeLU** (Lee et al. 2019). A meta-learning-based model built on the MAML framework, MeLU uses side information to quickly adapt to new tasks with minimal data, making it effective for cold-start scenarios.
- **Meta TL** (Wang, Ding, and Caverlee 2021). Meta TL addresses cold-start recommendations by modeling sequential user transitions through meta-learning, framing the problem as a few-shot learning task to enable rapid adaptation with limited interactions.

Experimental Results

Comparison Results. We evaluated the proposed CTM method against state-of-the-art baselines on three benchmark datasets: Electronics, MovieLens 100k, and KuaiRec. The results are presented in Table 2. The CTM method consistently outperforms all baselines across all datasets and metrics, demonstrating its effectiveness in cold-start scenarios. On the Electronics dataset, CTM notably outperforms the best baseline (Meta TL), with significant gains in MRR, NDCG@5, and Hit metrics. Similarly, on MovieLens 100k, CTM demonstrates superior performance, achieving substantial improvements over all baselines. CTM's performance on the KuaiRec dataset is especially remarkable, showing significant improvements across key metrics, particularly Hit@5 and Hit@1. These substantial gains highlight CTM's capacity to capture latent user preferences, driving superior recommendation accuracy in cold-start scenarios.

We conducted a Friedman test to evaluate the significance of performance differences among six methods across three datasets. At a 0.05 significance level, the critical F-test value is 4.103. The results in Table 3 indicate that all computed Friedman test statistic (τ_F) exceed the critical threshold, confirming significant performance differences among the algorithms.

Metrics	MRR	NDCG@5	Hits@5	Hits@1
τ_F	13.476	11.190	13.857	12.333

Table 3: Friedman Test Results

Ablation Study. We conducted an ablation study to evaluate the contribution of each module within the CTM model, validating their effectiveness and necessity. Two ablation experiments were designed: one excluding the Counterfactual Sequence Augmentation module (CSA), and the other excluding the Meta-path Neighbor Aggregation module (MNA). The experiments were performed on the KuaiRec and MovieLens 100k datasets, with results summarized in Table 4.

Without CSA : Model performance declined across both datasets when this module was removed, highlighting its importance in enhancing robustness through additional supervisory information. This module is especially valuable for improving adaptability and generalization in cold-start scenarios.

Without MNA : All metrics showed notable decreases across both datasets without this module, emphasizing its critical role in capturing latent item relationships and its essential contribution to user preference modeling.

The ablation study demonstrates that the Counterfactual Sequence Augmentation and Meta-path Neighbor Aggregation modules are essential for CTM's success, significantly enhancing user preference modeling and performance in cold-start tasks.

Dataset	Ablation Module	MRR	NDCG@5	Hits@5	Hits@1
KuaiRec	Complete Model	0.512	0.565	0.787	0.305
	Without CSA	0.476	0.499	0.734	0.248
	Without MNA	0.437	0.514	0.710	0.274
MovieLens 100k	Complete Model	0.347	0.363	0.539	0.220
	Without CSA	0.304	0.303	0.474	0.146
	Without MNA	0.312	0.317	0.521	0.154

Table 4: Ablation Study Results on the KuaiRec and MovieLens 100k Datasets

Parameter Analysis

- **Counterfactual Training Intervention Degree.** This analysis examines the impact of the counterfactual training intervention degree, denoted by σ, on model performance using the MovieLens 100k and KuaiRec datasets. As shown in Figure 4, σ significantly affects performance, particularly in datasets with higher interaction volumes, such as KuaiRec. Both datasets achieve optimal results when σ is set to 50%. Specifically, on the KuaiRec dataset, Hits@5 and NDCG@5 peak at $\sigma = 0.5$, while MRR and NDCG@1 are best at lower σ values. In contrast, the MovieLens 100k dataset shows more stable performance across σ variations, with lower σ values yielding improved MRR and NDCG@1.
- **Initial Number of Interactions.** We investigate the influence of the initial number of interactions, denoted by k, on model performance in cold-start scenarios. Figure 5 illustrates the comparative results across various key metrics for $k = 2$, $k = 3$, and $k = 4$. It reveals that the model consistently achieves superior performance when $k = 3$, outperforming the other configurations. This observation highlights the critical role of selecting an optimal initial number of interactions in enhancing the model's generalization ability and predictive accuracy, particularly in cold-start sequential recommendation tasks.

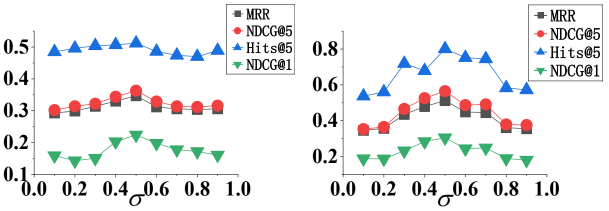

(a) MovieLens 100k (b) KuaiRec

Figure 4: Effect of σ on Model Performance in MovieLens 100k and Electronics

Conclusion and Future Work

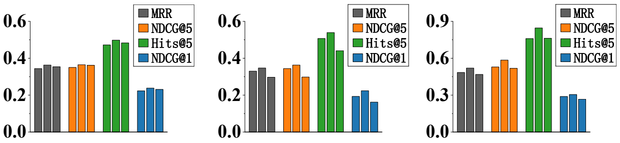

(a) Electronics (b) MovieLens 100k (c) KuaiRec

Figure 5: Effect of Initial Interactions k on Model Performance across Datasets: Electronics, MovieLens 100k, and KuaiRec

This paper introduces a novel cold-start sequential recommendation method that integrates meta-learning with counterfactual task augmentation. By generating counterfactual sequences through intervention in users' interaction histories, the proposed approach effectively captures potential shifts in user preferences resulting from different item interactions. The incorporation of meta-path neighbor aggregation enhances the model's ability to uncover latent relationships between items, enabling more detailed and accurate modeling of user preferences. Moreover, the meta-learning optimization component jointly optimizes both real and counterfactual task losses through a combination of global and local updates, ensuring robust adaptation in cold-start scenarios. Experimental evaluations on three publicly available datasets demonstrate that the proposed CTM method consistently outperforms five state-of-the-art cold-start recommendation methods, highlighting its effectiveness in modeling user preferences and their dynamic evolution. Future work will explore the impact of randomness in counterfactual sequence generation, aiming to enhance cold-start recommendation performance by examining how different sample selections affect meta-learning training.

Acknowledgments

This work is supported by the National Science and Technology Major Project (2020AAA0106102), the National Natural Science Foundation of China (62272285, U21A20473, 62072293, 62406180), the Shanxi Province Basic Research Program Project (202403021211086, 202203021211331), and the Key Technologies Program of Taihang Laboratory in Shanxi Province, China (THYF-JSZX-24010600).

References

Cai, D.; Qian, S.; Fang, Q.; Hu, J.; and Xu, C. 2023. User Cold-start Recommendation via Inductive Heterogeneous Graph Neural Network. *ACM Transactions on Information Systems*, 41(3): 1–27.

Cheng, H.-T.; Koc, L.; Harmsen, J.; Shaked, T.; Chandra, T.; Aradhye, H.; Anderson, G.; Corrado, G.; Chai, W.; Ispir, M.;

et al. 2016. Wide & Deep Learning for Recommender Systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Rystems*, 7–10.

Du, J.; Ye, Z.; Yao, L.; Guo, B.; and Yu, Z. 2022. Socially-aware Dual Contrastive Learning for Cold-start Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1927–1932.

Du, Y.; Zhu, X.; Chen, L.; Fang, Z.; and Gao, Y. 2023. MetaKG: Meta-learning on Knowledge Graph for Cold-start Recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(10): 9850–9863.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic Meta-learning for Fast Adaptation of Deep Networks. In *Proceedings of the International Conference on Machine Learning*, 1126–1135.

Gao, C.; Li, S.; Lei, W.; Chen, J.; Li, B.; Jiang, P.; He, X.; Mao, J.; and Chua, T.-S. 2022. KuaiRec: A Fully-observed Dataset and Insights for Evaluating Recommender Systems. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management*, 540–550.

Harper, F. M.; and Konstan, J. A. 2015. The Movielens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems*, 5(4): 1–19.

He, R.; Kang, W.-C.; and McAuley, J. 2017. Translation-based Recommendation. In *Proceedings of the 11th ACM Conference on Recommender Systems*, 161–169.

He, R.; and McAuley, J. 2016. Ups and downs: Modeling the Visual Evolution of Fashion Trends with One-class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web*, 507–517.

Huang, X.; Sang, J.; Yu, J.; and Xu, C. 2022. Learning to Learn a Cold-start Sequential Recommender. *ACM Transactions on Information Systems*, 40(2): 1–25.

Jin, B.; Gao, C.; He, X.; Jin, D.; and Li, Y. 2020. Multi-behavior Recommendation with Graph Convolutional Networks. In *Proceedings of the 43rd international ACM SIGIR Conference on Research and Development in Information Retrieval*, 659–668.

Jing, M.; Zhu, Y.; Zang, T.; and Wang, K. 2024. Contrastive Self-supervised Learning in Recommender Systems: A Survey. *ACM Transactions on Information Systems*, 42(2).

Kang, W.-C.; and McAuley, J. 2018. Self-attentive Sequential Recommendation. In *Proceedings of the 2018 IEEE International Conference on Data Mining*, 197–206.

Lee, H.; Im, J.; Jang, S.; Cho, H.; and Chung, S. 2019. Melu: Meta-learned User Preference Estimator for Cold-start Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1073–1082.

Man, T.; Shen, H.; Jin, X.; and Cheng, X. 2017. Cross-domain Recommendation: An Embedding and Mapping Approach. In *Proceedings of the International Joint Conferences on Artificial Intelligence*, volume 17, 2464–2470.

Peng, X.; Song, J.; and Liu, W. 2023. CALRec: Counterfactual Active Learning for Sequential Recommendation. In *Proceedings of the 6th International Conference on Artificial Intelligence and Big Data*, 186–191. IEEE.

Qian, T.; Liang, Y.; Li, Q.; and Xiong, H. 2023. Attribute Graph Neural Networks for Strict Cold Start Recommendation. In *Proceedings of the IEEE International Conference on Data Engineering*, 3783–3784.

Shu, H.; Chung, F.-L.; and Lin, D. 2023. Metagc-mc: A Graph-based Meta-learning Approach to Cold-start Recommendation with/without Auxiliary Information. *Information Sciences*, 623: 791–811.

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.

Volkovs, M.; Yu, G.; and Poutanen, T. 2017. Dropoutnet: Addressing Cold Start in Recommender Systems. In *Advances in Neural Information Processing Systems*, volume 30.

Wang, C.; Zhu, Y.; Sun, A.; Wang, Z.; and Wang, K. 2023. A Preference Learning Decoupling Framework for User Cold-Start Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1168–1177.

Wang, H.; Zhang, F.; Zhang, M.; Leskovec, J.; Zhao, M.; Li, W.; and Wang, Z. 2019. Knowledge-aware Graph neural Networks with Label Smoothness Regularization for Recommender Systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 968–977.

Wang, J.; Ding, K.; and Caverlee, J. 2021. Sequential Recommendation for Cold-start Users with Meta Transitional Learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1783–1787.

Wang, Z.; Zhang, J.; Xu, H.; Chen, X.; Zhang, Y.; Zhao, W. X.; and Wen, J.-R. 2021. Counterfactual Data-augmented Sequential Recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 347–356.

Wu, H.; Long, J.; Li, N.; Yu, D.; and Ng, M. K. 2022. Adversarial Auto-encoder Domain Adaptation for Cold-start Recommendation with Positive and Negative Hypergraphs. *ACM Transactions on Information Systems*, 41(2): 1–25.

Wu, S.; Wang, Y.; Jing, Q.; Dong, D.; Dou, D.; and Yao, Q. 2023. Coldnas: Search to Modulate for User Cold-start Recommendation. In *Proceedings of the ACM Web Conference*, 1021–1031.

Xu, Y.; Zhu, L.; Cheng, Z.; Li, J.; Zhang, Z.; and Zhang, H. 2023. Multi-Modal Discrete Collaborative Filtering for Efficient Cold-Start Recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(1): 741–755.

Zheng, Y.; Liu, S.; Li, Z.; and Wu, S. 2021. Cold-start Sequential Recommendation via Meta Learner. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4706–4713.

Zhu, Y.; Ge, K.; Zhuang, F.; Xie, R.; Xi, D.; Zhang, X.; Lin, L.; and He, Q. 2021a. Transfer-meta Framework for Cross-domain Recommendation to Cold-start Users. In *Proceedings of the 44th international ACM SIGIR Conference on Research and Development in Information Retrieval*, 1813–1817.

Zhu, Y.; Xie, R.; Zhuang, F.; Ge, K.; Sun, Y.; Zhang, X.; Lin, L.; and Cao, J. 2021b. Learning to Warm-up Cold Item Embeddings for Cold-start Recommendation with Meta Scaling and Shifting Networks. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1167–1176.