

Learning From Atypical Behavior: Temporary Interest Aware Recommendation Based on Reinforcement Learning

Ziwen Du, Ning Yang^{ID}, Member, IEEE, Zhonghua Yu, and Philip S. Yu^{ID}, Fellow, IEEE

Abstract—Traditional robust recommendation methods view atypical user-item interactions as noise and aim to reduce their impact with some kind of noise filtering technique, which often suffers from two challenges. First, in real world, atypical interactions may signal users' temporary interest different from their general preference. Therefore, simply filtering out the atypical interactions as noise may be inappropriate and degrade the personalization of recommendations. Second, it is hard to acquire the temporary interest since there are no explicit supervision signals to indicate whether an interaction is atypical or not. To address this challenges, we propose a novel model called Temporary Interest Aware Recommendation (TIARec), which can distinguish atypical interactions from normal ones without supervision and capture the temporary interest as well as the general preference of users. Particularly, we propose a reinforcement learning framework containing a recommender agent and an auxiliary classifier agent, which are jointly trained with the objective of maximizing the cumulative return of the recommendations made by the recommender agent. During the joint training process, the classifier agent can judge whether the interaction with an item recommended by the recommender agent is atypical, and the knowledge about learning temporary interest from atypical interactions can be transferred to the recommender agent, which makes the recommender agent able to alone make recommendations that balance the general preference and temporary interest of users. At last, the experiments conducted on real world datasets verify the effectiveness of TIARec.

Index Terms—Robust recommendation, temporary interest, deep reinforcement learning

1 INTRODUCTION

IN recent years, recommender systems have received much attention due to their ability to improve user experience in online applications like Amazon, Netflix, TikTok, etc. By nature, a recommender system makes recommendations based on the preference learned from users' historical interactions with items. However, in reality there always exist atypical behaviors that seldom appear before and obviously deviate from the general preference revealed by users' usual behaviors. For example, a user who is interested in football may occasionally watch some basketball-related videos during the finals of NBA, and when the NBA season is over, her attention to football will come back. The traditional works often consider the atypical interactions as noise that will impede the learning of the true preference of users and consequently weaken the robustness of recommender systems [3].

Recently, researchers have proposed a few models for robust recommendation, whose main idea is to reduce the

impact of atypical interactions with some kind of noise filtering technique [19], [20], [22], [24], [26]. Simply filtering out atypical interactions, however, may be inappropriate for recommender systems in real world. On the one hand, due to the dynamic nature of user preference, even so called atypical interactions do not have to mean nothing but can also reveal users' temporary interest. Consider the above NBA example again. Although the behaviors of watching basketball-related videos seldom happen before, they can still be a signal of the user's temporary interest in basketball due to NBA. Therefore, simply removing atypical interactions as noise may degrade the learning of user preference, especially the learning of user temporary interest. On the other hand, it is hard for the existing recommendation models to acquire the temporary interest directly from users' historical interactions, since there are no explicit supervision signals to indicate atypical interactions [25].

To further illustrate our motivation, we conduct an exploration of three real datasets Movielens-1M, Beauty, and Tianchi, where the interactions are labeled in terms of the item category. Fig. 1a shows the proportions of the users who have at least one rare interaction label with frequency less than 5% of the average, while Fig. 1b shows the ratios of the users whose at least one rare interaction is predicted correctly by the robust recommendation models CDAE [22] and APR [8] in addition to the proposed model TIARec on each dataset. We can see that although there are more than 70% users with atypical behaviors in each dataset, the temporary preferences are discovered for no more than 10% of them by the existing two robust recommendation models CDAE and APR. Therefore, we need a new method that is

- Ziwen Du, Ning Yang, and Zhonghua Yu are with the School of Computer Science, Sichuan University, Chengdu 610017, China. E-mail: ddzzww123456@gmail.com, {yangning, yuzhonghua}@scu.edu.cn.
- Philip S. Yu is with the Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60607 USA. E-mail: psyu@uic.edu.

Manuscript received 30 Mar. 2021; revised 29 Sept. 2021; accepted 15 Jan. 2022.

Date of publication 21 Jan. 2022; date of current version 15 Sept. 2023.

This work was supported in part by the National Natural Science Foundation of China under Grant 61972270, and in part by the NSF under Grants III-1763325, III-1909323, III-2106758, and SaTC-1930941.

(Corresponding author: Ning Yang.)

Recommended for acceptance by L. Chen, X. Zhou, X. Yang, and T. Sellis.

Digital Object Identifier no. 10.1109/TKDE.2022.3144292




Fig. 1. (a) Statistics of the users with rare interactions in real datasets Movielens-1M, Beauty, and Tianchi. (b) The recalls of the users with atypical interactions provided by the two state-of-the-art robust recommendation models CDAE, APR, and the proposed model TIARec.

able to distinguish whether an interaction is atypical or not and capture the temporary interest from atypical interactions.

To address the above challenges, motivated by the success of the reinforcement learning (RL) based recommendation models [5], [7], [9], [15], [17], [31], [35], we propose a novel RL-based model, called Temporary Interest Aware Recommendation (TIARec), which is able to distinguish atypical interactions from normal ones without annotations and capture the temporary interest as well as the general preference. TIARec contains a recommender agent and an auxiliary classifier agent. The recommender agent recommends an item to a user as its action at each time step, while the classifier agent can help to identify whether the interaction with the recommended item is atypical in a trial-error fashion. We also introduce a critic network and employ an Actor-Critic algorithm [13] to jointly train the recommender agent and the classifier agent with an objective to maximize the expected cumulative return of the recommendations made by the recommender agent. During the joint training, the knowledge about how to capture the temporary interest can be transferred from the classifier agent to the recommender agent. Once the training is finished, the recommender agent can alone make recommendations that balance the general preference and temporary interest of users.

Our contributions can be summarized as follows:

- We propose a novel RL-based model for robust recommendation, called Temporary Interest Aware Recommendation (TIARec). To our best knowledge, TIARec is the first reinforcement learning based solution to capture temporary interests for robust recommendation.
- We propose a deep reinforcement learning network including a recommender agent and an auxiliary classifier agent. During their joint training, the auxiliary classifier agent teaches the recommender agent the knowledge about learning the temporary interest from atypical interactions.
- Extensive experiments conducted on real world datasets verify the superiority of TIARec over the state-of-the-art methods on robust recommendation and RL-based recommendation.

The rest of this paper is organized as follows. In Section 2, we introduce the preliminaries and formally define the target problem. In Section 3, we first give an overview of TIARec and then describe its details. We empirically evaluate the performance of TIARec over real-world datasets in Section 4. At last, we briefly review the related works in Section 5 and conclude in Section 6.

2 PRELIMINARIES

In this section, we first give some basic definitions, then formulate the recommendations as a Markov Decision Process (MDP), and formally define the target problem of this paper.

2.1 Basic Definitions and Vectorization

Let \mathcal{U} be the user set and \mathcal{V} be the item set. Let $\mathcal{O} = \{o_1, \dots, o_{|\mathcal{O}|}\}$ be an interaction sequence, where $o_i \in \mathcal{V}$, $1 \leq i \leq |\mathcal{O}|$, is the item of the i th interaction. Let \mathcal{O}_t , \mathcal{M}_t , and \mathcal{N}_t be the sets of all interactions, normal interactions, and atypical interactions of some user that happen before time t , respectively. For the same user, $\mathcal{M}_t \cap \mathcal{N}_t = \Phi$ and $\mathcal{O}_t = \mathcal{M}_t \cup \mathcal{N}_t$. For vectorizing, we pretrain a d -dimensional vector for each item using Skip-gram algorithm, where an item o_i is treated as a word and a historical interaction sequence \mathcal{O}_t as a sentence. For simplicity, in the rest of this paper we use the pre-trained embeddings to represent the corresponding items, which are denoted by bold lowercase, for example o_i , and reuse \mathcal{V} to denote the whole set of the item embeddings.

2.2 MDP Formulation

Now we formulate the recommendations as an MDP, which is formally defined as follows:

- *State space* \mathcal{S} In recommender systems, the environment of the agents are defined as the users. Therefore the environment state is defined by a user's interaction history. Particularly, a state s is defined as any possible interaction set \mathcal{O} , i.e., the state space $\mathcal{S} = \{s = \mathcal{O}\}$.
- *Action space* \mathcal{A} An action a of a recommender system is defined as an embedding of a recommended item, and therefore the action space $\mathcal{A} = \mathcal{V}$. In this paper we assume only one item will be recommended at each time step.
- *Transitions* \mathcal{P} The transition $p(s'|s, a)$ is the conditional probability that the state of the environment (user) changes from $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ after an action $a \in \mathcal{A}$ is issued. In recommender systems, $p(s'|s, a)$ is 1 if the user interacts with the recommended item, otherwise 0.
- *Reward* \mathcal{R} For a recommender, the immediate reward $r(s, a)$ after issuing action a at state s will be defined in terms of users' feedbacks.
- *Discount rate* γ $\gamma \in (0, 1]$ is the discount factor for measuring long-term rewards.

2.3 Problem Formulation

Let $\pi : \mathcal{S} \times \mathcal{V} \mapsto [0, 1]$ be a recommender policy that outputs the conditional probability $p(a|s)$ of recommending item $a \in \mathcal{V}$ at state $s \in \mathcal{S}$. Let $\mathcal{V}^u = \{v_1^u, \dots, v_T^u\}$ be a sequence of the collected historical interactions of some user u that really happen before time step T , where T is the number of time steps considered and $v_i^u \in \mathcal{V}$ ($1 \leq i \leq T$) is the item embedding of the i th interaction. Given the dataset $\{\mathcal{V}^u\}$ over all users, we want to learn an optimal policy π^* for the recommender agent that maximizes the expected cumulative reward of the recommender, i.e.,

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right], \quad (1)$$

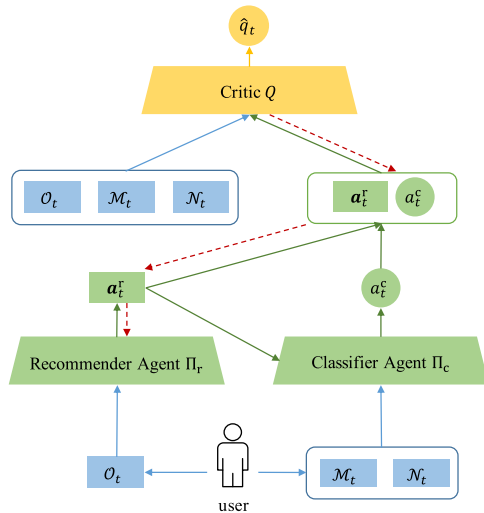


Fig. 2. The architecture of TIARec.

where $s_t = \mathcal{O}_t \in \mathcal{S}$ and $a_t \in \mathcal{V}$ are the user state and the recommended item at time step t , respectively.

3 THE PROPOSED MODEL

In this section, we first give an overview of the architecture of TIARec, and then describe its components and learning in detail.

3.1 Architecture of TIARec

As shown in Fig. 2, TIARec consists of three parts, a recommender agent Π_r , an auxiliary classifier agent Π_c , and a critic Q . The goal of TIARec is to learn an optimal recommender agent Π_r that realizes the objective defined by Equation (1), with the help of the classifier agent and the critic.

The agents Π_r and Π_c are two policy networks. At each time step t , based on the current state $s_t^r = \mathcal{O}_t$, the recommender agent recommends an item

$$a_t^r = \Pi_r(s_t^r; \Theta_r) \in \mathcal{V}, \quad (2)$$

as the action, where $a_t^r \in \mathbb{R}^{d \times 1}$ is an item embedding and Θ_r represents the trainable parameters of Π_r . Then the classifier agent judges how likely the interaction with a_t^r is atypical. The state s_t^c observed by the classifier agent includes three parts, the action a_t^r of the recommender agent, the current normal interaction set \mathcal{M}_t and atypical interaction set \mathcal{N}_t . Based on its current state $s_t^c = \{a_t^r, \mathcal{M}_t, \mathcal{N}_t\}$, the classifier agent outputs a scalar

$$a_t^c = \Pi_c(s_t^c; \Theta_c) \in [0, 1], \quad (3)$$

where Θ_c is the trainable parameters of Π_c . The real number a_t^c is a probability defining the action of the classifier agent, which is to recognize the interaction with a_t^r as atypical and append it to the atypical interaction set \mathcal{N}_t with probability a_t^c , or as a normal interaction and append it to the normal interaction set \mathcal{M}_t with probability $1 - a_t^c$. Essentially, \mathcal{N}_t encodes a user's temporary interest while \mathcal{M}_t encodes the user's general preference.

The critic network Q plays the role of the state-action value function in Actor-Critic algorithm, which helps the recommender agent and the classifier agent understand

how good the states and the actions are in terms of the expected cumulative return. With the overall state $s_t^q = \{\mathcal{O}_t, \mathcal{M}_t, \mathcal{N}_t\}$ and the overall action $a_t^q = a_t^r \oplus a_t^c$ as input, where \oplus represents concat operation and $a_t^q \in \mathbb{R}^{(d+1) \times 1}$, the critic evaluates the expected cumulative return \hat{q}_t for both the recommender agent and the classifier agent,

$$\hat{q}_t = Q(s_t^q, a_t^q; \Theta_q), \quad (4)$$

where Θ_q represents the trainable parameters of the critic. The overall objective of the training is to learn the optimal Θ_r , Θ_c , and Θ_q that maximize the expected cumulative return \hat{q}_t .

The insight here is that the classifier agent Π_c can help capture the temporary interest carried by the atypical interactions it recognizes, due to the trial-error nature of reinforcement learning. As the two agents are jointly trained, the knowledge about learning temporary interest from atypical interactions can be transferred from the classifier agent to the recommender agent along the back propagation path marked by the red dashed lines in Fig. 2. For example, if the classifier agent correctly recognizes an interaction as atypical (or mistakes it as a normal one), then the cumulative return will consequently increase (or reduce), and during the joint training, the supervision signal offered by the critic network will not only guide the update of the classifier agent, but also enforce the parameters of the recommender agent to be adjusted such that the recommender agent can remember that in such case, it should recommend an appropriate item satisfying the user's temporary interest rather than her general preference, so as to maximize the cumulative return.

3.2 Recommender Agent Π_r

At each time step t , the task of the recommender agent $\Pi_r(s_t^r; \Theta_r)$ is to predict an item a_t^r that a given user will interact with next time according to the current state $s_t = \mathcal{O}_t$ of that user, which is implemented as an MLP combining an attention network. Particularly, we first generate the current preference embedding $z_t^o \in \mathbb{R}^{d \times 1}$ based on the user's current state with following attention network:

$$z_t^o = \sum_{i=1}^{|\mathcal{O}_t|} \alpha_i^o o_i, \quad (5)$$

$$\alpha_i^o = \frac{\exp(\mathbf{q}^T \sigma(\mathbf{W}^o o_i + \mathbf{b}^o))}{\sum_{j=1}^{|\mathcal{O}_t|} \exp(\mathbf{q}^T \sigma(\mathbf{W}^o o_j + \mathbf{b}^o))},$$

where $o_i \in \mathbb{R}^{d \times 1}$ represents the pretrained embedding of the i th item $o_i \in \mathcal{O}_t$, $\alpha_i^o \in \mathbb{R}$ is the attention coefficient, and $\mathbf{q} \in \mathbb{R}^{d \times 1}$, $\mathbf{W}^o \in \mathbb{R}^{d \times d}$, and $\mathbf{b}^o \in \mathbb{R}^{d \times 1}$ are the trainable query vector, transformation matrix, and bias terms, respectively, and $\sigma(\cdot)$ is a non-linear activation function.

Then we feed z_t^o into an MLP network to generate a prototype item embedding $a_t^p \in \mathbb{R}^{d \times 1}$,

$$a_t^p = \text{MLP}_r(z_t^o; \theta_r), \quad (6)$$

where θ_r represents the trainable parameters of MLP_r . Now the set of the trainable parameters of the recommender agent is $\Theta_r = \{\theta_r, \mathbf{W}^o, \mathbf{b}^o, \mathbf{q}\}$.

Note that a_t^p is an embedding of the *prototype item* that may not exist in \mathcal{V} but defines how the embedding of the best recommended item should look. To generate a real item as

the action and set its immediate reward, we generate a top- k list \mathcal{V}_k including the first k closest item embeddings of \mathbf{a}_t^p in terms of cosine similarity. The item whose embedding is closest to \mathbf{a}_t^p is chosen as the action of the recommender agent during the training,

$$\mathbf{a}_t^r = \operatorname{argmax}_{v \in \mathcal{V}} \cos(\mathbf{a}_t^p, v). \quad (7)$$

where $\cos(\cdot, \cdot)$ is the cosine similarity and $v \in \mathbb{R}^{d \times 1}$ is the pretrained embedding of item v .

For offline training, we need to obtain the immediate reward R_t^r of the action \mathbf{a}_t^r at time step t . The challenge here is that the logged data $\{\mathcal{V}^u\}$ were generated by an uncontrollable recommender system whose recommending policy is unknown, or in other words, the historical data are user feedbacks to another unknown agent different from ours. Therefore, during the offline training, we cannot directly set R_t^r with respect to whether or not the recommended list \mathcal{V}_k covers the logged item that the user actually interacts with at time step t . The existing works often apply off-policy strategy to fulfill the offline training, by which the target recommending policy can be updated according to the data generated by a different but known policy [2], [30]. However, the traditional offline training strategy is inapplicable to our case, since in our case, the policy generating the logged data is unknown. Inspired by the techniques used in [1], [4], [11], [36], our idea to overcome this dilemma is to mimic the online interactions of the recommender agent with the user by inferring how likely the user will interact with an item in the recommended list based on the logged offline data. In particular, before the training, we first learn a user latent factor $\mathbf{u} \in \mathbb{R}^{d \times 1}$ for each user $u \in \mathcal{U}$ and an item latent factor $\mathbf{v} \in \mathbb{R}^{d \times 1}$ for each item $v \in \mathcal{V}$, by performing probabilistic matrix factorization (PMF) over the whole dataset in advance. Then the probability of user u will interact with item v can be calculated as

$$p(\mathbf{u}, \mathbf{v}) = \sigma(\mathbf{u}^T \mathbf{v}), \quad (8)$$

where $\sigma(\cdot)$ is sigmoid function. Next, during the training, we set the immediate reward R_t^r to the inferred probability that user u will interact with at least one item in the recommended list \mathcal{V}_k , i.e.,

$$R_t^r = \frac{1}{k} \sum_{v \in \mathcal{V}_k} p(\mathbf{u}, \mathbf{v}). \quad (9)$$

The insight here is that the larger the probability that user will interact with the recommended list, the higher the quality of it, and the more reward the agent deserves.

3.3 Classifier Agent Π_c

The task of the classifier agent $\Pi_c(s_t^c = \{\mathbf{a}_t^r, \mathcal{M}_t, \mathcal{N}_t\}; \Theta_c)$ is to determine how likely an interaction with the item \mathbf{a}_t^r recommended by the recommender agent is atypical, based on current $\{\mathcal{M}_t, \mathcal{N}_t\}$. Similar to the structure of the recommender agent, the classifier agent is also implemented as an MLP but combining two attention networks.

At each time step t , to capture the temporary interest from the atypical interactions, the classifier agent first generates a temporary interest embedding $\mathbf{z}_t^n \in \mathbb{R}^{d \times 1}$ based on

the current atypical interaction set \mathcal{N}_t , with the following attention network:

$$\begin{aligned} \mathbf{z}_t^n &= \sum_{i=1}^{|\mathcal{N}_t|} \alpha_i^n \mathbf{n}_i, \\ \alpha_i^n &= \frac{\exp[(\mathbf{a}_t^r)^T \sigma(\mathbf{W}^n \mathbf{n}_i + \mathbf{b}^n)]}{\sum_{j=1}^{|\mathcal{N}_t|} \exp[(\mathbf{a}_t^r)^T \sigma(\mathbf{W}^n \mathbf{n}_j + \mathbf{b}^n)]}, \end{aligned} \quad (10)$$

where $\mathbf{n}_i \in \mathbb{R}^{d \times 1}$ represents the pretrained embedding of the i th item in \mathcal{N}_t , $\alpha_i^n \in \mathbb{R}$ is the attention coefficient, $\mathbf{W}^n \in \mathbb{R}^{d \times d}$ and $\mathbf{b}^n \in \mathbb{R}^{d \times 1}$ are the learnable transformation matrix and bias terms, respectively, and \mathbf{a}_t^r serves as the query vector. Similarly, the classifier agent also generate a general preference embedding $\mathbf{z}_t^m \in \mathbb{R}^{d \times 1}$ based on the current normal interaction set \mathcal{M}_t ,

$$\begin{aligned} \mathbf{z}_t^m &= \sum_{i=1}^{|\mathcal{M}_t|} \alpha_i^m \mathbf{m}_i, \\ \alpha_i^m &= \frac{\exp[(\mathbf{a}_t^r)^T \sigma(\mathbf{W}^m \mathbf{m}_i + \mathbf{b}^m)]}{\sum_{j=1}^{|\mathcal{M}_t|} \exp[(\mathbf{a}_t^r)^T \sigma(\mathbf{W}^m \mathbf{m}_j + \mathbf{b}^m)]}. \end{aligned} \quad (11)$$

Finally, the classifier agent feeds the concatenation of the embeddings \mathbf{z}_t^n and \mathbf{z}_t^m into an MLP to generate the probability a_t^c that \mathbf{a}_t^r is an atypical interaction,

$$a_t^c = \text{MLP}_c(\mathbf{z}_t^n \oplus \mathbf{z}_t^m; \theta_c) \in [0, 1], \quad (12)$$

where \oplus is the concatenation operation and θ_c represents the trainable parameters of MLP_c . Now we can see that the trainable parameter set of the classifier agent is $\Theta_c = \{\theta_c, \mathbf{W}^n, \mathbf{W}^m, \mathbf{b}^n, \mathbf{b}^m\}$.

As we mentioned before, the probability a_t^c actually defines two alternative actions for the classifier agent, where one is to recognize \mathbf{a}_t^r as atypical with probability a_t^c and append it to \mathcal{N}_t , i.e.,

$$\mathcal{N}_{t+1} = \mathcal{N}_t \cup \{\mathbf{a}_t^r\}, \quad (13)$$

and the other one is to recognize \mathbf{a}_t^r as a normal interaction with probability $1 - a_t^c$ and append it to \mathcal{M}_t , i.e.,

$$\mathcal{M}_{t+1} = \mathcal{M}_t \cup \{\mathbf{a}_t^r\}. \quad (14)$$

As we will see later, at the beginning of a training cycle, \mathcal{N}_0 and \mathcal{M}_0 are initialized with empty set, and they will gradually grow with more actions issued by the recommender agent.

The classifier agent also needs to obtain an immediate reward for its action. Intuitively, atypical behaviors should be dissimilar to each other and lead to diversity, otherwise many similar atypical behaviors will not be atypical any more. Therefore, if \mathbf{a}_t^r is recognized as atypical appropriately, the overall similarity between the items in \mathcal{N}_t should reduce after \mathbf{a}_t^r is merged into \mathcal{N}_t , which implies that smaller similarity between \mathbf{a}_t^r and the items in \mathcal{N}_t should receive greater reward. On the contrary, if \mathbf{a}_t^r is recognized as normal appropriately, the overall similarity between the items in \mathcal{M}_t should rise after \mathbf{a}_t^r is merged into \mathcal{M}_t , which implies that larger similarity between \mathbf{a}_t^r and the items in \mathcal{M}_t is better. Based on this idea, we set the immediate reward R_t^c for the classifier agent as

$$R_t^c = \begin{cases} 1 - \frac{1}{|\mathcal{N}_t|} \sum_i^{|\mathcal{N}_t|} \cos(\mathbf{n}_i, \mathbf{a}_t^r), & \text{if } a_t^r \in \mathcal{N}_{t+1}, \\ \frac{1}{|\mathcal{M}_t|} \sum_i^{|\mathcal{M}_t|} \cos(\mathbf{m}_i, \mathbf{a}_t^r), & \text{if } a_t^r \in \mathcal{M}_{t+1}. \end{cases} \quad (15)$$

There are two points here we want to further clarify. First, as there are no supervision signals, how can the classifier agent learn to recognize atypical behavior? Our idea here is that atypical behaviors cause dissimilarity, and the immediate reward given by Equation (15) and the expected cumulative return estimated by the critic (see next subsection) will judge how well the classifier agent classifies a recommended item as atypical or normal, and accordingly adjust its parameters through the policy gradient descent during the training process. Second, where is the temporary interest? As users' temporary interest hides in their atypical interactions, we argue that the atypical interaction set \mathcal{N}_t reveals what the classifier agent considers the temporary interest, and during the joint training, such knowledge will be transferred to the recommender agent.

3.4 Critic Q

The critic network $Q(s_t^q, \mathbf{a}_t^q; \Theta_q)$ is responsible for evaluating the expected cumulative return \hat{q}_t based on the state $s_t^q = \{\mathcal{O}_t, \mathcal{M}_t, \mathcal{N}_t\}$ and action $\mathbf{a}_t^q = \mathbf{a}_t^r \oplus a_t^c$. The critic generates the estimation of the expected cumulative return through an MLP,

$$\hat{q}_t = \text{MLP}_q(\mathbf{a}_t^q \oplus \mathbf{z}_t^o \oplus \mathbf{z}_t^n \oplus \mathbf{z}_t^m; \Theta_q), \quad (16)$$

where \mathbf{z}_t^o , \mathbf{z}_t^n , and \mathbf{z}_t^m are the preference embeddings of \mathcal{O}_t , \mathcal{M}_t , and \mathcal{N}_t , respectively, which are produced by Equations (5), (10), and (11), respectively.

3.5 Learning of TIARec

As we have mentioned, the recommender agent is trained jointly with the classifier agent and the critic network using an actor-critic algorithm. In particular, we employ MADDPG [13] algorithm to learn the parameters Θ_r , Θ_c , and Θ_q of TIARec.

3.5.1 Target Networks

The target networks are used to provide the supervision signal for the training. The target networks $\Pi_r'(s_{t+1}^r; \Theta_r')$ and $\Pi_c'(s_{t+1}^c; \Theta_c')$ are a copy of Π_r and Π_c , with their own parameters Θ_r' and Θ_c' , respectively, which are used to generate the next actions \mathbf{a}_{t+1}^r and a_{t+1}^c . Similarly, the target network $Q'(s_{t+1}^q, \mathbf{a}_{t+1}^q; \Theta_q')$ of the critic Q is a copy of Q with parameters Θ_q' , which is used for generating the target cumulative return when training Q .

3.5.2 Loss Function of The Critic

We expect the estimated cumulative return \hat{q}_t provided by the critic Q should be as close to the target cumulative return q_t as possible. Let r_t be the total immediate reward and $\{s_t^q, \mathbf{a}_t^q, r_t, s_{t+1}^q\}$ be a training sample, then the loss function of the critic network can be defined as

$$L_q(\Theta_q) = \mathbb{E}_{s_t, \mathbf{a}_t^q, r_t, s_{t+1}^q} [q_t - Q(s_t^q, \mathbf{a}_t^q = \mathbf{a}_t^r \oplus a_t^c; \Theta_q)]. \quad (17)$$

The total immediate reward r_t is defined as

$$r_t = R_t^r + \alpha R_t^c, \quad (18)$$

where R_t^r and R_t^c are the immediate rewards of the recommender agent and the classifier agent, respectively, and α is

a hyper-parameter for balancing the contributions of the immediate rewards.

The target cumulative return q_t is defined as the sum of the immediate reward and the cumulative return of $\{s_{t+1}^q, \mathbf{a}_{t+1}^q\}$,

$$q_t = r_t + \gamma Q'(s_{t+1}^q, \mathbf{a}_{t+1}^q; \Theta_q'), \quad (19)$$

where $\gamma \in [0, 1]$ is the discount factor. The next action $\mathbf{a}_{t+1}^q = \mathbf{a}_{t+1}^r \oplus a_{t+1}^c$, where $\mathbf{a}_{t+1}^r = \Pi_r'(s_{t+1}^r; \Theta_r')$ and $a_{t+1}^c = \Pi_c'(s_{t+1}^c; \Theta_c')$. As we will see later, in each iteration of the training, the parameters Θ_q' , Θ_r' , and Θ_c' of the target networks are fixed with the model parameters Θ_q , Θ_r , and Θ_c of previous iteration.

Algorithm 1. Training of TIARec

input: The training dataset $\{\mathcal{V}^u\}$, learning rate λ , fusion coefficient $\tau \in [0, 1]$

output: The parameters Θ_r , Θ_c , and Θ_q

- 1: Randomly initialize Θ_r , Θ_c , and Θ_q ;
- 2: Initialize the target networks: $\Theta_r' = \Theta_r$, $\Theta_c' = \Theta_c$, $\Theta_q' = \Theta_q$;
- 3: Initialize the capacity of replay buffer \mathcal{B} ;
- 4: **while** not convergent **do**
- 5: **foreach** \mathcal{V}^u **do**
- 6: Initialize $\mathcal{O}_0 = \mathcal{M}_0 = \mathcal{N}_0 = \{\}$;
- 7: **for** $t \in \{0, 1, \dots, T\}$ **do**
- 8: Generate the prototype item \mathbf{a}_t^p and the recommended item \mathbf{a}_t^r using Equations (6) and (7), respectively;
- 9: Set recommender reward R_t^r using Equation (9);
- 10: Update $\mathcal{O}_{t+1} = \mathcal{O}_t \cup \{\mathbf{a}_t^r\}$;
- 11: Generate classifier action a_t^c using Equation (12);
- 12: Update $\mathcal{N}_{t+1} = \mathcal{N}_t \cup \{\mathbf{a}_t^r\}$, $\mathcal{M}_{t+1} = \mathcal{M}_t$ with probability a_t^c ; or $\mathcal{N}_{t+1} = \mathcal{N}_t$, $\mathcal{M}_{t+1} = \mathcal{M}_t \cup \{\mathbf{a}_t^r\}$ with probability $1 - a_t^c$;
- 13: Set classifier reward R_t^c using Equation (15);
- 14: Set overall reward r_t using Equation (18);
- 15: Set $s_t = \{\mathcal{O}_t, \mathcal{N}_t, \mathcal{M}_t\}$ and $s_{t+1} = \{\mathcal{O}_{t+1}, \mathcal{N}_{t+1}, \mathcal{M}_{t+1}\}$;
- 16: Save transition sample $(s_t, \mathbf{a}_t^r, r_t, s_{t+1})$ to \mathcal{B} ;
- 17: Sample a minibatch from \mathcal{B} ;
- 18: Compute the losses $L_q, L_{r,c}$ based on the minibatch, using Equations (17) and (20), respectively;
- 19: Update the critic $\Theta_q = \Theta_q - \lambda \nabla_{\Theta_q} L_q$;
- 20: Update the recommender $\Theta_r = \Theta_r - \lambda \nabla_{\Theta_r} L_{r,c}$;
- 21: Update the classifier $\Theta_c = \Theta_c - \lambda \nabla_{\Theta_c} L_{r,c}$;
- 22: Update the target networks: $\Theta_q' = \tau \Theta_q + (1 - \tau) \Theta_q'$; $\Theta_r' = \tau \Theta_r + (1 - \tau) \Theta_r'$; $\Theta_c' = \tau \Theta_c + (1 - \tau) \Theta_c'$;
- 23: **end**
- 24: **end**
- 25: **end**

3.5.3 Loss Function of The Agents

As we have seen before, the recommender agent actually generates a prototype item embedding \mathbf{a}_t^p , while a real item embedding \mathbf{a}_t^r is fed to the critic. To connect the recommender agent and the critic for joint training, we need to make them as similar as possible, and at the same time, both agents seek to maximize the expected cumulative return \hat{q}_t . These ideas lead to the following loss function of the agents

$$L_{r,c}(\Theta_r, \Theta_c) = \mathbb{E}_{s_t} \left[(\mathbf{a}_t^r - \mathbf{a}_t^p)^2 - Q(s_t^q, \mathbf{a}_t^q = \mathbf{a}_t^p \oplus a_t^c) \right]. \quad (20)$$

Note that since the first term ensures the recommended item embedding \mathbf{a}_t^r will be close to the prototype embedding \mathbf{a}_t^p .

a_t^p , we use the prototype item a_t^p instead of a_t^r to compute the expected cumulative return, which is slightly different from Equation (17) where we use a_t^r .

3.5.4 Training Algorithm

The training of TIARec uses an offline strategy shown in Algorithm 1, where each iteration is composed of two stages, the sample generating and the policy updating. In the sample generating stage of each time step t (lines 8-16), the recommender agent first recommends an item a_t^r (line 8), and the classifier agent generates a probability a_t^c (line 11) with respect to which the interaction with the recommended item is classified as noise or normal (line 12). Then we set the overall immediate reward r_t (line 14) by fusing the immediate rewards R_t^r (line 9) and R_t^c (line 13) of the recommender agent and the classifier agent. At last, the state is updated (line 15) and the transition sample is stored to the replay buffer (line 16). In the policy updating stage (lines 17-22), we first sample a minibatch from the replay buffer (line 17), and then use gradient descent to update the parameters of the critic, the recommender agent, the classifier agent, and the target networks (lines 19-22).

4 EXPERIMENTS

The goals of the experiments are to answer the following three research questions:

- *RQ1* how does TIARec perform as compared to the state-of-the-art baselines?
- *RQ2* how does the classifier agent contribute to the performance of TIARec?
- *RQ3* how can the superiority of TIARec be illustrated with an intuitive and visualizable case study?
- *RQ4* how is the robustness of TIARec?
- *RQ5* how do the hyper-parameters affect the performance?

The implementation of TIARec is available on <https://github.com/dududu123456/TIARec>.

4.1 Experimental Setup

4.1.1 Datasets

We adopt the following three public datasets for the experiments.

- *Movielens-1M*¹ Movielens-1M is a popular dataset for evaluating recommendation models, which contains over 1 million ratings given by 6,040 users to 3,952 movies.
- *Beauty*² Beauty contains over 2 million ratings given by more than 1 million users to 249,274 products on Amazon.
- *Tianchi*³ Tianchi contains over 802,757 ratings given by 6,471 users to 554,442 products on TMALL, a well-known Chinese e-commerce website.

TABLE 1
Statistics of the Datasets

	Movielens-1M	Beauty	Tianchi
The number of users	6,040	1,210,271	6,471
The number of items	3,952	249,274	554,442
The number of ratings	1,000,209	2,023,096	802,757

The statistics of the datasets is given in Table 1. On each dataset, we use the first 70% of the data as the training set, the middle 10% as the validation set, and the remaining 20% as the testing set.

4.1.2 Baselines

As we have mentioned, TIARec is a RL-based robust recommendation model that can capture users' temporary interest. Therefore, to verify our claimed contributions, we compare TIARec with seven representative methods including two robust recommendation models (CDAE and APR), four RL-based models (DEERs, HRL, KGQR, and PDQ), and one interactive recommendation model (I-CARS), which are briefly introduced as follows:

- *CDAE* CDAE [22] is a Denoising Auto-encoder based collaborative filtering model, which can learn robust latent representations of corrupted user-item interactions for top- k recommendation.
- *APR* APR [8] is an Adversarial Personalized Ranking framework, which offers the recommendation robustness by injecting adversarial noise to the parameters of a pairwise ranking model BPR [16] during an adversarial training.
- *DEERs* DEERs [31] is a deep reinforcement learning based recommendation model, which can simultaneously model positive and negative feedbacks using a Deep Q-network (DQN) [15] with a pairwise training.
- *HRL* HRL [25] is a hierarchical reinforcement learning model for recommendations of online courses, which directly remove atypical interactions through a hierarchical sequential decision process.
- *KGQR* KGQR [34] is a knowledge graph enhanced Q-learning model for interactive recommendations, which incorporates the prior knowledge of the item correlation learned from knowledge graph into the sequential decision making to improve sample efficiency for learning better representations of items and users.
- *PDQ* PDQ [36] is another deep Q-learning based interactive recommendation model, where the recommendation policy is trained without the requirement of real customer interactions, but with a pseudo environment model that simulates the feedbacks of users.
- *I-CARS* I-CARS [14] is an interactive context-aware recommendation model, which can continually update its parameters according to a user's recent preferences that are revealed by the user's feedbacks to the deliberately selected items.

Additionally, in order to verify the effectiveness of the classifier agent of TIARec, we also compare TIARec with

1. <https://grouplens.org/datasets/movielens/>
 2. <https://www.kaggle.com/skillsmuggler/amazon-ratings>
 3. <https://tianchi.aliyun.com/dataset/dataDetail?dataId=64345>

TABLE 2
Performance Comparison on Movielens-1M

Method	Metric	HR@k			Recall@k			NDCG@k		
		k=5	k=10	k=20	k=5	k=10	k=20	k=5	k=10	k=20
CDAE		0.0554	0.0561	0.3014	0.0134	0.0297	0.0674	0.0003	0.0044	0.0187
APR		0.0536	0.0613	0.3387	0.0121	0.0315	0.0752	0.0034	0.0151	0.0304
DEERs		0.0815	0.1166	0.3651	0.0259	0.0477	0.1134	0.0083	0.0174	0.0381
HRL		0.0734	0.1284	0.3433	0.0299	0.0311	0.0998	0.0087	0.0188	0.0377
KGQR		0.1391	0.2765	0.5029	0.0283	0.0519	0.1378	0.0118	0.0165	0.0545
PDQ		0.0829	0.1619	0.3712	0.0174	0.0418	0.1049	0.0126	0.0196	0.0364
I-CARS		0.0629	0.0929	0.3406	0.0150	0.0356	0.1009	0.0025	0.0089	0.0219
TIARec-C		0.1104	0.1916	0.4128	0.0231	0.0505	0.1043	0.0072	0.0170	0.0325
TIARec		0.1706	0.3110	0.5793	0.0322	0.0583	0.1602	0.0151	0.0274	0.0701
Improv.		22.65%	12.48%	15.19%	7.69%	12.33%	12.26%	19.84%	39.80%	28.62%

TABLE 3
Performance Comparison on Beauty

Method	Metric	HR@k			Recall@k			NDCG@k		
		k=5	k=10	k=20	k=5	k=10	k=20	k=5	k=10	k=20
CDAE		0.0024	0.0936	0.2431	0.0009	0.0161	0.0339	0.0017	0.0054	0.0151
APR		0.0085	0.1077	0.3668	0.0018	0.0137	0.0548	0.0018	0.0081	0.0216
DEERs		0.1025	0.1902	0.3066	0.0054	0.0217	0.0514	0.0041	0.0077	0.0184
HRL		0.1033	0.1178	0.2433	0.0105	0.0233	0.0531	0.0049	0.0118	0.0197
KGQR		0.1120	0.1714	0.3731	0.0114	0.0183	0.0425	0.0056	0.0152	0.0166
PDQ		0.1328	0.2074	0.4378	0.0160	0.0212	0.0548	0.0073	0.0124	0.0183
I-CARS		0.1044	0.1704	0.3728	0.0104	0.0177	0.0357	0.0021	0.0071	0.0163
TIARec-C		0.1009	0.2027	0.3101	0.0151	0.0229	0.0409	0.0071	0.0130	0.0173
TIARec		0.1652	0.2439	0.4699	0.0231	0.0328	0.0701	0.0104	0.0207	0.0304
Improv.		24.40%	17.60%	7.33%	44.38%	40.77%	27.92%	42.47%	36.18%	40.74%

one more baseline method TIARec-C, which is a variant of TIARec where the classifier agent is removed.

4.1.3 Evaluation Protocol

We choose the widely used Hit Rate (HR), Recall, and NDCG to evaluate the performance of TIARec and the baselines. Let $\{v_1^u, \dots, v_{N_u}^u\}$ be the testing set consisting of N_u ground-truth items that user u has interacted with. For each testing instance v_i^u ($1 \leq i \leq N_u$), we generate a ranking list l_i^u of items using TIARec. Then the number of u 's cases that a ground-truth is ranked in the top- k items can be computed as

$$S_k^u = \sum_{i=1}^{N_u} \mathbb{I}(\text{rank}(v_i^u, l_i^u) \leq k), \quad (21)$$

where $\text{rank}(v, l)$ is the rank of item v in list l , and $\mathbb{I}(x)$ is the indicator function that returns 1 if x is true, otherwise 0. Then the metrics can be defined as follows:

$$\text{HR@}k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbb{I}(S_k^u \geq 1), \quad (22)$$

$$\text{Recall@}k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{S_k^u}{N_u}, \quad (23)$$

$$\text{NDCG@}k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{N_u} \sum_{i=1}^{N_u} \frac{\mathbb{I}(\text{rank}(v_i^u, l_i^u) \leq k)}{\log_2(1 + \text{rank}(v_i^u, l_i^u))}. \quad (24)$$

4.1.4 Hyper-Parameter Setting

The hyper-parameters are tuned on the validation sets. The embedding dimensionality d is set to 64 for Movielens-1M and 128 for Beauty and Tianchi. The reward balance factor

α in Equation (18) is set to 0.1 for Movielens-1M, 0.5 for Beauty and 1.0 for Tianchi. On three datasets, the discount rate γ , the size of replay buffer \mathcal{B} , and the fusion coefficient τ are set to 0.99, 2,000 and 0.01, respectively. For fairness, the hyper-parameters of the baselines are also set to their optimal values on the validation sets.

4.2 Comparative Study (RQ1)

Tables 2, 3, and 4 show the results of the comparison between TIARec and the baselines on Movielens-1M, Beauty, and Tianchi, respectively, where the last row shows the percentage by which TIARec improves the performance of the best baseline.

At first, we can note that on both datasets TIARec outperforms the robust recommendation models. In contrast to the robust models CDAE and APR, which rely on some kind of noise filtering technique, TIARec can learn from the atypical interactions rather than filtering them. Particularly, TIARec can learn the knowledge about how to capture the users' temporary interest from the atypical interactions, and transfer it to the recommender agent from the classifier agent during the joint training, which makes TIARec more powerful than the traditional robust recommendation models.

At the same time, we can observe that TIARec achieves significant superiority over the RL-based baselines, which demonstrates the advantages of the classifier agent of TIARec. In particular, unlike the traditional RL-based models, which offers no robustness since they indiscriminately treat every user-item interaction as normal, TIARec can distinguish atypical interactions from normal ones with the help of the classifier agent during the joint training with the recommender agent, which brings two advantages. First, the differential treatment reduces the impact of the atypical

TABLE 4
Performance Comparison on Tianchi

Method	Metirc			HR@k			Recall@k			NDCG@k		
	k=5	k=10	k=20	k=5	k=10	k=20	k=5	k=10	k=20	k=5	k=10	k=20
CDAE	0.0034	0.0172	0.1669	0.0011	0.0073	0.0218	0.0003	0.0017	0.0083			
APR	0.0079	0.0745	0.1967	0.0028	0.0109	0.0332	0.0007	0.0053	0.0088			
DEERs	0.0363	0.0818	0.2446	0.0046	0.0102	0.0366	0.0014	0.0037	0.0114			
HRL	0.0847	0.0940	0.2327	0.0092	0.0107	0.0344	0.0039	0.0090	0.0117			
KGQR	0.0863	0.1417	0.2632	0.0105	0.0165	0.0405	0.0034	0.0069	0.0131			
PDQ	0.0771	0.1255	0.2517	0.0085	0.0120	0.0378	0.0027	0.0036	0.0121			
I-CARS	0.0238	0.0515	0.1675	0.0006	0.0060	0.0122	0.0003	0.0029	0.0091			
TIARec-C	0.0526	0.1037	0.2199	0.0082	0.0160	0.0392	0.0022	0.0074	0.0112			
TIARec	0.0922	0.1515	0.2713	0.0119	0.0194	0.0433	0.0062	0.0104	0.0154			
Improv.	8.85%	6.92%	3.08%	13.33%	17.58%	6.91%	58.97%	15.56%	17.56%			

interactions for learning users' general preference. Second, the joint training implicitly teaches the recommender agent how to capture users' temporary interest. Such two advantages enable the recommender agent of TIARec to make recommendations with a balance between users' general preference and temporary interest, which contrasts sharply with the traditional RL-based models.

We can also see that TIARec performs better than I-CARS on each dataset. I-CARS aims at capturing a user's preference shift according to the user's feedbacks to the well-chosen probing questions, which is similar to the target of TIARec. However, there are two points that tell I-CARS apart from TIARec. First, I-CARS considers each feedback to be a signal of a user's true preference, which makes I-CARS more susceptible to the noise incurred by the user's feedbacks to the probing questions which cannot reflect the user's true preference. In contrast, the recommender agent of TIARec can learn to discern whether an atypical behavior represents the user's true temporary interest or just a noise, by the joint training with the auxiliary classifier agent. Second, essentially I-CARS is a greedy algorithm which wants to fit every observation, while TIARec pursues the maximization of the long-term return due to the nature of reinforcement learning. For the same reason, almost in each case all the RL-based models achieve a better performance than I-CARS. However, we can see that I-CARS still outperforms CDAE and APR in most cases, due to its sensitivity to user's preference changes.

At last, Fig. 3 illustrates the training process of TIARec in terms of the average cumulative return estimated by the critic over all the users, which is defined as

$$\text{Average-Q} = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{1}{T} \sum_{t=0}^T \hat{q}_t^u, \quad (25)$$

where \hat{q}_t^u is the expected cumulative return estimated by the critic for user u at time step t (see Equation (16)). From Fig. 3




Fig. 3. Training procedure on all datasets.

we can see that as the number of training epochs increases, the Average-Q first improves, and then approaches convergence after 50 epochs on MovieLens-1M, 200 epochs on Beauty, and 30 epochs on Tianchi. Note that it is reasonable that the training on Beauty takes longer time before convergence since it has much greater scale of data.

4.3 Ablation Study (RQ2)

Now we investigate the effectiveness of the classifier agent of TIARec. At first we compare TIARec with its variant TIARec-C where the classifier agent is removed. From Tables 2, 3, and 4 we can observe that TIARec considerably improves the performance of the variant TIARec-C on three datasets, which demonstrates the indispensable role of the classifier agent for the robustness of TIARec. As we have mentioned before, the classifier agent can identify whether an interaction is atypical during the joint training with the recommender agent with the objective of maximizing the expected cumulative return estimated by the critic. In fact, once the classifier agent is removed, TIARec degenerates to a traditional model which loses the capability of learning users' temporary interest from atypical interactions.

The classifier agent plays its role during the training process through its immediate reward in Equation (18), where its weight is regularized by the coefficient α . To get further insight into the classifier agent, Fig. 4 illustrates how the performance of the recommender agent changes with the different weights α of the classifier agent on both datasets. At first, when $\alpha = 0$, the classifier agent contributes nothing and TIARec degenerates to TIARec-C. As α increases, the classifier agent accounts for more proportion of the total immediate reward and more knowledge about learning temporary interest from atypical interactions is transferred to the recommender agent, which leads to the better performance of the latter. However, when α is greater than 0.1 on MovieLens-1M, 0.5 on Beauty, and 1.0 on Tianchi, the curves begin dropping. This is because an overlage α results in the total immediate reward overwhelmed by the classifier agent and consequently causes the recommender agent to pay excessive attention to the users' temporary interest.

4.4 Case Study (RQ3)

Now we conduct a case study to illustrate TIARec's ability to capture temporary interest and make recommendations with balance between the general preference and temporary




Fig. 4. The effectiveness of the classifier agent.

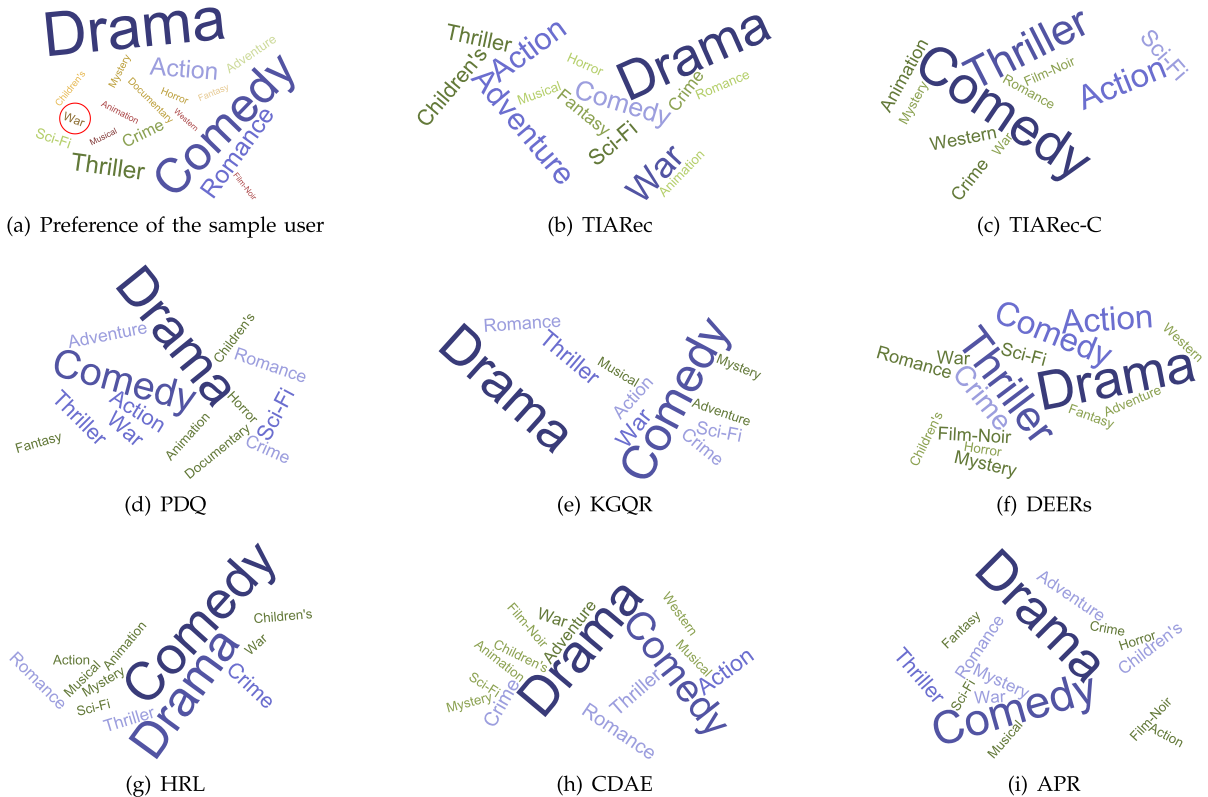


Fig. 5. Case study.

interest of a user. For this purpose, we sample one user with ID "264" from Movielens-1M. Fig. 5a uses a word cloud to visualize the tag distribution of the movies watched by the sample user in 2002, where the size of a tag word is proportional to the frequency that the sample user watched the movies with that tag. From Fig. 5a we can see that in 2002, the sample user generally preferred drama and comedy, and seldom watched movies of other types. However, in October and November the sample user watched some war movies for some unknown reason, as marked with a red circle in Fig. 5a. We randomly choose one war movie the sample user watched in November as the ground-truth for test, and check the top-20 movies recommended by TIARec and the baselines, whose tag distributions are shown in Figs. 5b, 5c, 5d, 5e, 5f, 5g, 5h, and 5i, respectively.

From Fig. 5b we can see that the drama movies and war movies account for the vast majority of the top-20 list recommended by TIARec. In fact, 6 war movies together with 10 drama movies are recommended in the top-20 list, including 4 movies of both drama and war, which illustrates that TIARec is able to capture the sample user's temporary

interest in war movie and can balance it with her general preference to drama when making recommendations. We also note that as another general preference, the comedy movies are much less than the drama ones. We argue that this is because as war movie is more compatible with drama movie, quite a few war movies are also tagged with drama. To make an easy tradeoff between the temporary interest and the general preference of the sample user in this case, TIARec tends to combine war movies with drama rather than comedy ones.

From Fig. 5c we can see that without the help of the auxiliary classifier agent, TIARec-C pays most attention to the sample user's general preference to comedy movies and can hardly recommend war movies. Compared with Figs. 5b and 5c shows that it is the classifier agent that offers TIARec the ability to learn a user's temporary interest from her atypical behaviors.

Figs. 5d, 5e, 5f, 5g, 5h, and 5i show that similar to TIARec-C, the sample user's general preference to drama or comedy dominates the recommendations made by the baselines. In fact, none of them recommends more than 2 war





Fig. 6. Robustness over different noise levels.

Fig. 7. Tuning of the embedding size d .

movies, which indicates that the baselines can hardly capture the sample user's temporary interest in war movie from the recent watching of the war movies, and excessively rely on the general preference of users when making recommendations. Compared with Fig. 5b, such results verify the superiority of TIARec in the temporal interest learning for robust recommendation.

4.5 Robustness Study (RQ4)

Now we investigate the robustness of TIARec by comparing it with the baseline methods on the three datasets, in terms of the metric HR@10. For this purpose, in each testing set, with respect to a given ratio (noise level) we randomly change some consecutive five interactions to the items of the same category whose frequency is less than the 20% of the average, to simulate the atypical behaviors that can be regarded as temporary interest. Fig. 6 shows the results at noise levels 0.1, 0.2, 0.4, and 0.8. From Fig. 6 we can see that on each dataset the performances of all methods decline as noise level rises. However, TIARec still consistently outperforms the baseline methods at different noise levels, which again shows the advantage of TIARec in capturing temporary interest of a user. In particular, it is worthy to note that when the noise level is 0.8, the interest of almost every user becomes very unstable and changes quite frequently (every five interactions), which results in a drastic deterioration of the performances of all the methods. However, even in such extreme case, the HR@10 of TIARec still remains greater than 0.1 which is remarkably higher than that of the baseline methods.

4.6 Hyper-Parameter Tuning (RQ5)

At last we tune the embedding size d on validation sets, of which the results are shown in Fig. 7. We can observe that the performance curves first increase and then decrease when we enlarge the embedding size d . The performance reaches its maximum at $d = 64$ on Movielens-1M and at $d = 128$ on both Beauty and Tianchi. A larger d means more

model parameters needed to be learned, which may lead to the problem of overfitting.

5 RELATED WORK

In this section, we briefly review two domains of the works mostly related to our study, including robust recommendation and RL-based recommendation.

5.1 Robust Recommendation

The goal of robust recommendation is to reduce the impact of atypical interactions on user preference learning as much as possible to improve the robustness of the models. The existing methods for robust recommendation can be roughly divided into two classes. One class of the methods employs auto-encoder based techniques to generate robust embeddings for users and items [10], [18], [22]. For example, Wu *et al.* propose the model CDAE [22] which first injects random drop-out noise to training data and trains a denoising auto-encoder [20] on intentionally corrupted input with the objective of minimizing reconstruction errors. Li *et al.* [10] propose a collaborative variational auto-encoder (CVAE) for robust recommendation of multimedia, and Shenbin *et al.* [18] propose a Recommender VAE (RecVAE) model with that can be trained with corrupted implicit user-item interaction vectors. The other class of the robust recommendation methods introduces adversarial noise as well as adversarial training to improve the model robustness [8], [19], [21], [24]. For example, He *et al.* propose the model APR [8] which enhances the pairwise ranking method BPR [16] with an adversarial training. Similarly, Wang *et al.* [21], Tang *et al.* [19], and Yuan *et al.* [24] propose to add perturbations to the input data and use adversarial learning to strengthen the robustness of recommendation model. As we have mentioned, the existing robust recommendations methods often neglect the meaningful signal carried by the atypical interactions, from which our TIARec can learn users' temporary interest for making better recommendations.

5.2 RL-Based Recommendation

Recently, RL-based recommendation approaches have attracted more attentions from scholars. RL-based models aim to learn an optimal strategy to maximize the long terms rewards. RL-based models can be divided into three categories, the policy-based methods [1], [2], [4], [12], [29], [30], [32], the value-based methods [11], [31], [33], [36], and the actor-critic based methods [5], [6], [7], [23], [27], [28]. Chen *et al.* [1] propose to use a balanced hierarchical clustering tree to tackle the large action space problem. Chen *et al.* [2] propose to use an off-policy framework to solve the data bias problem. Liu *et al.* [12] model the interactions between users and items to represent users' state. Zhao *et al.* [31] take the negative feedback and positive feedback into consideration to represent users' state. Zheng *et al.* [33] utilize two Q-network to balance the correlation between exploit and explore. Zou *et al.* [36] construct a simulator to model users' feedback to handle the selection bias of logged data. Zhou *et al.* [34] propose to obtain item representations through a graph convolution network and obtain users' state embedding vectors through a recurrent neural network, and then trains the recommendation strategy with the Q-learning algorithm. He *et al.* [7] and Feng *et al.* [5] treat each agent as a scenario and use a multi-agent framework to improve the total performance of all scenarios. Zhang *et al.* [27] apply a multi-agent reinforcement learning method to coauthor network analysis for the dynamic collaboration recommendation. Zhao *et al.* [32] propose an actor-critic based reinforcement learning method for list-wise recommendation and introduce a simulator of online user-agent interacting environment to get more interactions. Zhao *et al.* [29] also propose a multi-agent reinforcement learning model to capture the sequential correlation among different scenarios. The existing RL-based recommendation methods equally treat user-item interactions as normal, which offer little robustness against atypical interactions and cannot capture users' temporary interest too.

6 CONCLUSION

In this paper, we propose a novel reinforcement learning based model for robust recommendation, called Temporary Interest Aware Recommendation (TIARec), which can learn users' temporary interest from their atypical interactions. TIARec contains a recommender agent and an auxiliary classifier agent, which are jointly trained with the objective of maximizing the expected cumulative return of the recommendations made by the recommender agent. In particular, during the joint training, the classifier agent can judge whether the interaction with a recommended item is atypical and transfer the knowledge about the temporary interest learning to the recommender agent, which finally enable the recommender agent to alone make robust recommendations that balance the general preference and temporary interest of users. The experiments conducted on real world datasets show that due to the ability to capture users' temporary interest, TIARec achieves a significant improvement in recommendation performance compared with the state-of-the-art baselines.

REFERENCES

- [1] H. Chen *et al.*, "Large-scale interactive recommendation with tree-structured policy gradient," in *Proc. 33rd AAAI Conf. Artif. Intell.*, 2019, pp. 3312–3320.
- [2] M. Chen, A. Beutel, P. Covington, S. Jain, F. Belletti, and E. H. Chi, "Top-k off-policy correction for a reinforce recommender system," in *Proc. 12th ACM Int. Conf. Web Search Data Mining*, 2019, pp. 456–464.
- [3] Y. Deldjoo, T. Di Noia, E. Di Sciascio, and F. A. Merra, "How dataset characteristics affect the robustness of collaborative recommendation models," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 951–960.
- [4] G. Dulacarnold, R. Evans, P. Sunehag, and B. Coppin, "Deep reinforcement learning in large discrete action spaces," 2015, *arXiv:1512.07679*.
- [5] J. Feng *et al.*, "Learning to collaborate: Multi-scenario ranking via multi-agent reinforcement learning," in *Proc. World Wide Web Conf.*, 2018, pp. 1939–1948.
- [6] T. Gui *et al.*, "Mention recommendation in twitter with cooperative multi-agent reinforcement learning," in *Proc. 42nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2019, pp. 535–544.
- [7] X. He *et al.*, "Learning to collaborate in multi-module recommendation via multi-agent reinforcement learning without communication," in *Proc. 14th ACM Conf. Recommender Syst.*, 2020, pp. 210–219.
- [8] X. He, Z. He, X. Du, and T.-S. Chua, "Adversarial personalized ranking for recommendation," in *Proc. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2018, pp. 355–364.
- [9] E. Ie *et al.*, "SLATEQ: A tractable decomposition for reinforcement learning with recommendation sets," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, 2019, pp. 2592–2599.
- [10] X. Li and J. She, "Collaborative variational autoencoder for recommender systems," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 305–314.
- [11] F. Liu, H. Guo, X. Li, R. Tang, Y. Ye, and X. He, "End-to-end deep reinforcement learning based recommendation with supervised embedding," in *Proc. 13th Int. Conf. Web Search Data Mining*, 2020, pp. 384–392.
- [12] F. Liu *et al.*, "Deep reinforcement learning based recommendation with explicit user-item interactions modeling," 2018, *arXiv:1810.12027*.
- [13] R. Lowe, Y. Wu, A. Tamar, J. Harb, O. P. Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6382–6393.
- [14] R. Lumbantoran, X. Zhou, Y. Ren, and L. Chen, "I-CARS: An interactive context-aware recommender system," in *Proc. Int. Conf. Data Mining*, 2019, pp. 1240–1245.
- [15] V. Mnih *et al.*, "Playing Atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.
- [16] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "BPR: Bayesian personalized ranking from implicit feedback," in *Proc. 25th Conf. Uncertainty Artif. Intell.*, 2009, pp. 452–461.
- [17] W. Shang, Y. Yu, Q. Li, Z. Qin, Y. Meng, and J. Ye, "Environment reconstruction with hidden confounders for reinforcement learning based recommendation," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 566–576.
- [18] I. Shenbin, A. Alekseev, E. Tutubalina, V. Malykh, and S. I. Nikolenko, "RecVAE: A new variational autoencoder for top-N recommendations with implicit feedback," in *Proc. 13th Int. Conf. Web Search Data Mining*, 2020, pp. 528–536.
- [19] J. Tang, X. Du, X. He, F. Yuan, Q. Tian, and T. Chua, "Adversarial training towards robust multimedia recommender system," *IEEE Trans. Knowl. Data Eng.*, vol. 32, no. 5, pp. 855–867, May 2020.
- [20] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proc. 25th Int. Conf. Mach. Learn.*, 2008, pp. 1096–1103.
- [21] J. Wang *et al.*, "IRGAN: A minimax game for unifying generative and discriminative information retrieval models," in *Proc. 40th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2017, pp. 515–524.
- [22] Y. Wu, C. DuBois, A. X. Zheng, and M. Ester, "Collaborative denoising auto-encoders for top-N recommender systems," in *Proc. 9th ACM Int. Conf. Web Search Data Mining*, 2016, pp. 153–162.
- [23] X. Xin, A. Karatzoglou, I. Arapakis, and J. M. Jose, "Self-supervised reinforcement learning for recommender systems," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 931–940.


- [24] F. Yuan, L. Yao, and B. Benatallah, "Adversarial collaborative neural network for robust recommendation," in *Proc. 42nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2019, pp. 1065–1068.
- [25] J. Zhang, B. Hao, B. Chen, C. Li, H. Chen, and J. Sun, "Hierarchical reinforcement learning for course recommendation in MOOCs," in *Proc. 33rd AAAI Conf. Artif. Intell.*, 2019, pp. 435–442.
- [26] S. Zhang, H. Yin, T. Chen, Q. V. N. Hung, Z. Huang, and L. Cui, "GCN-based user representation learning for unifying robust recommendation and fraudster detection," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 689–698.
- [27] Y. Zhang, C. Zhang, and X. Liu, "Dynamic scholarly collaborator recommendation via competitive multi-agent reinforcement learning," in *Proc. 11th ACM Conf. Recommender Syst.*, 2017, pp. 331–335.
- [28] D. Zhao, L. Zhang, B. Zhang, L. Zheng, Y. Bao, and W. Yan, "MaHRL: Multi-goals abstraction based deep hierarchical reinforcement learning for recommendations," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 871–880.
- [29] X. Zhao, L. Xia, D. Yin, and J. Tang, "Model-based reinforcement learning for whole-chain recommendations," 2019, *arXiv:1902.03987*.
- [30] X. Zhao, L. Xia, L. Zhang, Z. Ding, D. Yin, and J. Tang, "Deep reinforcement learning for page-wise recommendations," in *Proc. 12th ACM Conf. Recommender Syst.*, 2018, pp. 95–103.
- [31] X. Zhao, L. Zhang, Z. Ding, L. Xia, J. Tang, and D. Yin, "Recommendations with negative feedback via pairwise deep reinforcement learning," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1040–1048.
- [32] X. Zhao, L. Zhang, L. Xia, Z. Ding, D. Yin, and J. Tang, "Deep reinforcement learning for list-wise recommendations," 2017, *arXiv:1801.00209*.
- [33] G. Zheng *et al.*, "DRN: A deep reinforcement learning framework for news recommendation," in *Proc. World Wide Web Conf.*, 2018, pp. 167–176.
- [34] S. Zhou *et al.*, "Interactive recommender system via knowledge graph-enhanced reinforcement learning," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 179–188.
- [35] L. Zou, L. Xia, Z. Ding, J. Song, W. Liu, and D. Yin, "Reinforcement learning to optimize long-term user engagement in recommender systems," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 2810–2818.
- [36] L. Zou *et al.*, "Pseudo Dyna-Q: A reinforcement learning framework for interactive recommendation," in *Proc. 13th Int. Conf. Web Search Data Mining*, 2020, pp. 816–824.




Ziwen Du received the bachelor's degree in 2018 from the School of Computer Science, Sichuan University, China, where he is currently working toward the master's degree with the School of Computer Science. His research interests include data mining and recommender systems.



Ning Yang (Member, IEEE) received the PhD degree in computer science from Sichuan University in 2010. He is currently an associate professor with Sichuan University, China. His research interests include recommender systems and social media mining.



Zhonghua Yu received the PhD degree in computer science from Belarusian State University in 1995. He is currently an associate professor with Sichuan University, China. From 2001 to 2003, he was a postdoctor with the University of Tokyo, under the supervision of Junichi Tsujii. His research interests include natural language processing and machine learning.



Philip S. Yu (Fellow, IEEE) received the PhD degree in electrical engineering from Stanford University. He is currently a distinguished professor of computer science with the University of Illinois at Chicago and the Wexler chair of information technology. His research interests include Big Data, data mining, and social computing. He is a fellow of the ACM.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.