Recommendations with Negative Feedback via Pairwise Deep Reinforcement Learning

Xiangyu Zhao
Data Science and Engineering Lab
Michigan State University
zhaoxi35@msu.edu

Liang Zhang
Data Science Lab
Intelligent Advertising Lab
JD.com
zhangliang16@jd.com

Zhuoye Ding
Data Science Lab
JD.com
dingzhuoye@jd.com

Long Xia
Data Science Lab
JD.com
xialong@jd.com

Jiliang Tang
Data Science and Engineering Lab
Michigan State University
tangjili@msu.edu

Dawei Yin
Data Science Lab
JD.com
yindawei@acm.org

ABSTRACT

Recommender systems play a crucial role in mitigating the problem of information overload by suggesting users' personalized items or services. The vast majority of traditional recommender systems consider the recommendation procedure as a static process and make recommendations following a fixed strategy. In this paper, we propose a novel recommender system with the capability of continuously improving its strategies during the interactions with users. We model the sequential interactions between users and a recommender system as a Markov Decision Process (MDP) and leverage Reinforcement Learning (RL) to automatically learn the optimal strategies via recommending trial-and-error items and receiving reinforcements of these items from users' feedback. Users' feedback can be positive and negative and both types of feedback have great potentials to boost recommendations. However, the number of negative feedback is much larger than that of positive one; thus incorporating them simultaneously is challenging since positive feedback could be buried by negative one. In this paper, we develop a novel approach to incorporate them into the proposed deep recommender system (DEERS) framework. The experimental results based on real-world e-commerce data demonstrate the effectiveness of the proposed framework. Further experiments have been conducted to understand the importance of both positive and negative feedback in recommendations.

KEYWORDS

Recommender System, Deep Reinforcement Learning, Pairwise Deep Q-Network.

ACM Reference Format:

Xiangyu Zhao, Liang Zhang, Zhuoye Ding, Long Xia, Jiliang Tang, and Dawei Yin. 2018. Recommendations with Negative Feedback via Pairwise Deep

Reinforcement Learning. In *KDD '18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August 19–23, 2018, London, United Kingdom*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3219819.3219886>

1 INTRODUCTION

Recommender systems are intelligent E-commerce applications. They assist users in their information-seeking tasks by suggesting items (products, services, or information) that best fit their needs and preferences. Recommender systems have become increasingly popular in recent years, and have been utilized in a variety of domains including movies, music, books, search queries, and social tags [23, 24]. Typically, a recommendation procedure can be modeled as interactions between users and recommender agent (RA). It consists of two phases: 1) user model construction and 2) recommendation generation [14]. During the interaction, the recommender agent builds a user model to learn users' preferences based on users' information or feedback. Then, the recommender agent generates a list of items that best match users' preferences.

Most existing recommender systems including collaborative filtering, content-based and learning-to-rank consider the recommendation procedure as a static process and make recommendations following a fixed greedy strategy. However, these approaches may fail given the dynamic nature of the users' preferences. Furthermore, the majority of existing recommender systems are designed to maximize the immediate (short-term) reward of recommendations, i.e., to make users purchase the recommended items, while completely overlooking whether these recommended items will lead to more profitable (long-term) rewards in the future [25].

In this paper, we consider the recommendation procedure as sequential interactions between users and recommender agent; and leverage Reinforcement Learning (RL) to automatically learn the optimal recommendation strategies. Recommender systems based on reinforcement learning have two advantages. First, they are able to continuously update their try-and-error strategies during the interactions, until the system converges to the optimal strategy that generates recommendations best fitting their users' dynamic preferences. Second, the models in the system are trained via estimating the present value with delayed rewards under current states

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '18, August 19–23, 2018, London, United Kingdom

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5552-0/18/08...\$15.00

<https://doi.org/10.1145/3219819.3219886>

Figure 1: Impact of negative feedback on recommendations.

and actions. The optimal strategy is made by maximizing the expected long-term cumulative reward from users. Therefore the system could identify items with small immediate rewards but making big contributions to the rewards for future recommendations.

Efforts have been made on utilizing reinforcement learning for recommender systems [25, 27]. For instance, the work [25] modeled the recommender system as a MDP process and estimated the transition probability and then the Q-value table. However, these methods may become inflexible with the increasing number of items for recommendations. This prevents them from being adopted in e-commerce recommender systems. Thus, we leverage Deep Q-Network (DQN), an (adapted) artificial neural network, as a non-linear approximator to estimate the action-value function in RL. This model-free reinforcement learning method does not estimate the transition probability and not store the Q-value table. This makes it flexible to support huge amount of items in recommender systems. It can also enrich the system’s generalization compared to traditional approaches that estimate action-value function separately for each sequence.

When we design recommender systems, positive feedback (clicked / ordered feedback) represents the users’ preference and thus is the most important information to make recommendations. In reality, users also skip some recommended items during the recommendation procedure. These skipped items influence user’s click/order behaviors [6], which can help us gain better understandings about users’ preferences. Hence, it is necessary to incorporate such negative feedback. However, the number of skipped items (or negative feedback) is typically far larger than that of positive ones. Hence, it is challenging to capture both positive and negative feedback since positive feedback could be buried by negative one. In this paper, we propose a framework DEERS to model positive and negative feedback simultaneously. As shown in Figure 1, when items are skipped by a user, traditional recommender systems do not update their strategies, and still recommend similar items; while the proposed framework can capture these negative feedback and update the strategy for recommendations. We summarize our major contributions as follows:

- We identify the importance of negative feedback from the recommendation procedure and provide a principled approach to capture it for recommendations;

- We propose a deep reinforcement learning based framework DEERS, which can automatically learn the optimal recommendation strategies by incorporating positive and negative feedback;
- We demonstrate the effectiveness of the proposed framework in real-world e-commerce data and validate the importance of negative feedback in accurate recommendations.

The rest of this paper is organized as follows. In Section 2, we formally define the problem of recommender system via reinforcement learning. In Section 3, we provide approaches to model the recommending procedure as sequential user-agent interactions and introduce details about employing deep reinforcement learning to automatically learn the optimal recommendation strategies via offline-training. Section 4 carries out experiments based on real-world offline users’ behavior log and presents experimental results. Section 5 briefly reviews related work. Finally, Section 6 concludes this work and discusses our future work.

2 PROBLEM STATEMENT

We study the recommendation task in which a recommender agent (RA) interacts with environment $\mathcal{E}$ (or users) by sequentially choosing recommendation items over a sequence of time steps, so as to maximize its cumulative reward. We model this problem as a Markov Decision Process (MDP), which includes a sequence of states, actions and rewards. More formally, MDP consists of a tuple of five elements $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ as follows:

- **State space $\mathcal{S}$:** A state $s_t \in \mathcal{S}$ is defined as the browsing history of a user before time t . The items in s_t are sorted in chronological order.
- **Action space $\mathcal{A}$:** The action $a_t \in \mathcal{A}$ of RA is to recommend items to a user at time t . Without the loss of generality, we assume that the agent only recommends one item to the user each time. Note that it is straightforward to extend it with recommending multiple items.
- **Reward $\mathcal{R}$:** After the RA taking an action a_t at the state s_t , i.e., recommending an item to a user, the user browses this item and provides her feedback. She can skip (not click), click, or order this item, and the RA receives immediate reward $r(s_t, a_t)$ according to the user’s feedback.
- **Transition probability $\mathcal{P}$:** Transition probability $p(s_{t+1}|s_t, a_t)$ defines the probability of state transition from s_t to s_{t+1} when RA takes action a_t . We assume that the MDP satisfies $p(s_{t+1}|s_t, a_t, \dots, s_1, a_1) = p(s_{t+1}|s_t, a_t)$.
- **Discount factor γ :** $\gamma \in [0, 1]$ defines the discount factor when we measure the present value of future reward. In particular, when $\gamma = 0$, RA only considers the immediate reward. In other words, when $\gamma = 1$, all future rewards can be counted fully into that of the current action.

Figure 2 illustrates the agent-user interactions in MDP. With the notations and definitions above, the problem of item recommendation can be formally defined as follows: *Given the historical MDP, i.e., $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, the goal is to find a recommendation policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$, which can maximize the cumulative reward for the recommender system.*

Figure 2: The agent-user interactions in MDP.

3 THE PROPOSED FRAMEWORK BASED ON DEEP REINFORCEMENT LEARNING WITH NEGATIVE FEEDBACKS

Negative feedback dominates users' feedback to items and positive feedback could be buried by negative one if we aim to capture them simultaneously. Thus in this section, we first propose a basic model that only considers positive feedback (clicked/ordered items) into the state. We build the deep reinforcement learning framework under this setting. After that, we consider negative feedback (skipped items) in the state space and redesign the deep architecture. Finally, we discuss how to train the framework via offline users' behavior log and how to utilize the framework for item recommendations.

3.1 The Basic DQN Model

The positive items represent the key information about users' preference, i.e., which items the users prefer to. A good recommender system should recommend the items users prefer the most. Thus, we only consider the positive feedback into the state in the basic model. More formally, we redefine the state and transition process for the basic model below:

- **State s :** State $s = \{i_1, \dots, i_N\} \in \mathcal{S}$ is defined as the previous N items that a user clicked/ordered recently. The items in state s are sorted in chronological order.
- **Transition from s to s' :** When RA recommends item a at state s to a user, the user may skip, click, or order the item. If the user skips the item, then the next state $s' = s$; while if the user clicks/orders the item, then the next state $s' = \{i_2, \dots, i_N, a\}$.

Note that in reality, using discrete indexes to denote items is not sufficient since we cannot know the relations between different items only from indexes. One common way is to use extra information to represent items. For instance, we can use the attribute information like brand, price, sale per month, etc. Instead of extra item information, our model uses the RA/user interaction information, i.e., users' browsing history. We treat each item as a word and the clicked items in one recommendation session as a sentence. Then, we can obtain dense and low-dimensional vector representations for items by word embedding[11].

As shown in Figure 2, by interacting with the environment (users), recommender agent takes actions (recommends items) to users in such a way that maximizes the expected return, which includes the delayed rewards. We follow the standard assumption that delayed rewards are discounted by a factor of γ per time-step, and define the action-value function $Q(s, a)$ as the expected return based on state s and action a . The optimal action-value function $Q^*(s, a)$,

Figure 3: The architecture of the basic DQN model for recommendations.

which has the maximum expected return achievable by the optimal policy, should follow the Bellman equation [2] as:

$$Q^*(s, a) = \mathbb{E}_{s'} [r + \gamma \max_{a'} Q^*(s', a') | s, a]. \quad (1)$$

In real recommender systems, the state and action spaces are enormous, thus estimating $Q^*(s, a)$ by using the Bellman equation for each state-action pair is infeasible. In addition, many state-action pairs may not appear in the real trace such that it is hard to update their values. Therefore, it is more flexible and practical to use an approximator function to estimate the action-value function, i.e., $Q^*(s, a) \approx Q(s, a; \theta)$. In practice, the action-value function is usually highly nonlinear. Deep neural networks are known as excellent approximators for non-linear functions. In this paper, We refer to a neural network function approximator with parameters θ as a Q-network. A Q-network can be trained by minimizing a sequence of loss functions $L(\theta)$ as:

$$L(\theta) = \mathbb{E}_{s, a, r, s'} [(y - Q(s, a; \theta))^2], \quad (2)$$

where $y = \mathbb{E}_{s'} [r + \gamma \max_{a'} Q(s', a'; \theta^P) | s, a]$ is the target for the current iteration. The parameters from the previous iteration θ^P are fixed when optimizing the loss function $L(\theta)$. The derivatives of loss function $L(\theta)$ with respect to parameters θ are presented as follows:

$$\begin{aligned} \nabla_{\theta} L(\theta) = \mathbb{E}_{s, a, r, s'} [(r + \gamma \max_{a'} Q(s', a'; \theta^P) \\ - Q(s, a; \theta)) \nabla_{\theta} Q(s, a; \theta)]. \end{aligned} \quad (3)$$

In practice, it is often computationally efficient to optimize the loss function by stochastic gradient descent, rather than computing the full expectations in the above gradient. Figure 3 illustrates the architecture of basic DQN model.

3.2 The Proposed DEERS Framework

As previously discussed, the items that users skipped can also indicate users' preferences, i.e., what users may not like. However, the system with only positive items will not change its state or update its strategy when users skip the recommended items. Thus, the state should not only contain positive items that user clicked

or ordered, but also incorporate negative (skipped) items. To integrate negative items into the model, we redefine state and transition process as follows:

- **State** $s: s = (s_+, s_-) \in S$ is a state, where $s_+ = \{i_1, \dots, i_N\}$ is defined as the previous N items that a user clicked or ordered recently, and $s_- = \{j_1, \dots, j_N\}$ is the previous N items that the user skipped recently. The items in s_+ and s_- are processed in chronological order.
- **Transition from s to s'** : When RA recommends item a at state $s = (s_+, s_-)$ to a user, if users skip the recommended item, we keep $s'_+ = s_+$ and update $s'_- = \{j_2, \dots, j_N, a\}$. If users click or order the recommended item, update $s'_+ = \{i_2, \dots, i_N, a\}$, while keeping $s'_- = s_-$. Then set $s' = (s'_+, s'_-)$.

We can see from the new definition that no matter users accept or skip recommendations, the new system will incorporate the feedback for next recommendations. Also it can distinguish positive and negative feedback with s_+ and s_- .

3.2.1 The Architecture of DEERS Framework. In this approach, we concatenate positive state s_+ and a recommended item a as positive input (positive signals), while concatenating negative state s_- and the recommended item a as the negative input (negative signals). Rather than following the basic reinforcement learning architecture in Figure 3, we construct DQN in a novel way. Figure 4 illustrates our new DQN architecture. Instead of just concatenating clicked/ordered items $\{i_1, \dots, i_N\}$ as positive state s_+ , we introduce a RNN with Gated Recurrent Units (GRU) to capture users' sequential preference as positive state s_+ :

$$z_n = \sigma(W_z i_n + U_z h_{n-1}), \quad (4)$$

$$r_n = \sigma(W_r i_n + U_r h_{n-1}), \quad (5)$$

$$\hat{h}_n = \tanh[W i_n + U(r_n \cdot h_{n-1})], \quad (6)$$

$$h_n = (1 - z_n)h_{n-1} + z_n \hat{h}_n, \quad (7)$$

where GRU utilizes an update gate z_n to generate a new state and a reset gate r_n to control the input of the former state h_{n-1} . The inputs of GRU are the embeddings of user's recently clicked/ordered items $\{i_1, \dots, i_N\}$, while we use the output (final hidden state h_N) as the representation of the positive state, i.e., $s_+ = h_N$. We obtain negative state s_- following a similar way. Here we leverage GRU rather than Long Short-Term Memory (LSTM) because that GRU outperforms LSTM for capturing users' sequential behaviors in recommendation task [8].

As shown in Figure 4, we feed positive input (positive signals) and negative input (negative signals) separately in the input layer. Also, different from traditional fully connected layers, we separate the first few hidden layers for positive input and negative input. The intuition behind this architecture is to recommend an item that is similar to the clicked/ordered items (left part), while dissimilar to the skipped items (right part). This architecture can assist DQN to capture distinct contributions of the positive and negative feedback to recommendations.

Then, the loss function $L(\theta)$ for an iteration of the new deep Q-network training can be rewritten as:

$$L(\theta) = \mathbb{E}_{s, a, r, s'} \left[\left(y - Q(s_+, s_-, a; \theta) \right)^2 \right], \quad (8)$$

Figure 4: The architecture of the DEERS framework.

where $y = \mathbb{E}_{s'}[r + \gamma \max_{a'} Q(s'_+, s'_-, a'; \theta^p) | s_+, s_-, a]$ is the target of current iteration. Then, the gradient of loss function with respect to the θ can be computed as:

$$\nabla_{\theta} L(\theta) = \mathbb{E}_{s, a, r, s'} \left[\left(r + \gamma \max_{a'} Q(s'_+, s'_-, a'; \theta^p) - Q(s_+, s_-, a; \theta) \right) \nabla_{\theta} Q(s_+, s_-, a; \theta) \right]. \quad (9)$$

To decrease the computational cost of Equation (9), we adopt item recalling mechanism to reduce the number of relevant items¹.

3.2.2 The Pairwise Regularization Term. With deep investigations on the users' logs, we found that in most recommendation sessions, the RA recommends some items belong to the same category (e.g. telephone), while users click/order a part of them and skip others. We illustrate a real example of a recommendation session in Table 1, in which three categories of items are recommended to user, and each time the RA recommends one item to the user. For category B, we can observe that the user clicked item a_2 while skipped item a_5 , which indicates the partial order of user's preference over these two items in category B. This partial order naturally inspires us maximizing the difference of Q-values between $Q(s_2, a_2)$ and $Q(s_5, a_5)$. At time 2, we name a_5 as the *competitor item* of a_2 . Sometimes, one item could have multiple "competitor" items; thus we select the item at the closest time as the "competitor" item. For example, at time 3, a_3 's competitor item is a_1 rather than a_6 . Overall, we select one target item's *competitor item* according to three requirements: 1) the "competitor" item belongs to the same category with the target item; 2) user gives

¹In general, user's current preference should be related to user's previous clicked/ordered items (say $\mathcal{L}_p$). Thus for each item in $\mathcal{L}_p$, we collect a number of most similar items from the whole item space, and combine all collected items as the item space for current recommendation session.

Table 1: An illustrative example of a recommendation session.

Time	State	Item	Category	Feedback
1	s_1	a_1	A	skip
2	s_2	a_2	B	click
3	s_3	a_3	A	click
4	s_4	a_4	C	skip
5	s_5	a_5	B	skip
6	s_6	a_6	A	skip
7	s_7	a_7	C	order

different types of feedback to the “competitor” item and the target item; and 3) the “competitor” item is at the closest time to the target item.

To maximize the difference of Q-values between target and competitor items, we add a regularization term to Equation 8:

$$L(\theta) = \mathbb{E}_{s, a, r, s'} \left[\left(y - Q(s_+, s_-, a; \theta) \right)^2 - \alpha \left(Q(s_+, s_-, a; \theta) - Q(s_+, s_-, a^C; \theta) \right)^2 \right], \quad (10)$$

where $y = \mathbb{E}_{s'} [r + \gamma \max_{a'} Q(s'_+, s'_-, a'; \theta^P) | s_+, s_-, a]$ is the target for the iteration. The second term aims to maximize the difference of Q-values between target item a and competitor item a^C at state (s_+, s_-) , which is controlled by a non-negative parameter α . Note that since user’s preference is relatively stable during a short time slot [30], we assume that user will give same feedback to a^C at state (s_+, s_-) . For example, if RA recommends item a_5 at state s_2 , the user will still skip a_5 . The gradient of loss function can be computed as:

$$\begin{aligned} \nabla_{\theta} L(\theta) = \mathbb{E}_{s, a, r, s'} \left[\left(y - Q(s_+, s_-, a; \theta) \right) \nabla_{\theta} Q(s_+, s_-, a; \theta) \right. \\ \left. - \alpha \left(Q(s_+, s_-, a; \theta) - Q(s_+, s_-, a^C; \theta) \right) \right. \\ \left. \left(\nabla_{\theta} Q(s_+, s_-, a; \theta) - \nabla_{\theta} Q(s_+, s_-, a^C; \theta) \right) \right]. \end{aligned} \quad (11)$$

3.3 Off-policy Training Task

With the proposed deep reinforcement learning framework, we can train the parameters in models and then do the test work. We train the proposed model based on users’ offline log, which records the interaction history between RA’s policy $b(s_t)$ and users’ feedback. RA takes the action based on the off-policy $b(s_t)$ and obtain the feedback from the offline log. We present our off-policy training algorithm in details shown in Algorithm 1. Note that our algorithm is *model-free*: it solves the reinforcement learning task directly using samples from the environment, without explicitly constructing an estimate of the environment.

In each iteration on a training recommendation session, there are two stages. For storing transitions stage: given the state s_t , the recommender agent first recommends an item a_t from a fixed off-policy $b(s_t)$ (line 7), which follows a standard off-policy way

Algorithm 1 Off-policy Training of DEERS Framework.

```

1: Initialize the capacity of replay memory  $\mathcal{D}$ 
2: Initialize action-value function  $Q$  with random weights
3: for session = 1,  $M$  do
4:   Initialize state  $s_0$  from previous sessions
5:   for  $t = 1, T$  do
6:     Observe state  $s_t = \{i_1, \dots, i_N, j_1, \dots, j_N\}$ 
7:     Execute action  $a_t$  following off-policy  $b(s_t)$ 
8:     Observe reward  $r_t$  from users
9:     Set  $s_{t+1} = \begin{cases} \{i_2, \dots, i_N, a_t, j_1, \dots, j_N\} & r_t > 0 \\ \{i_1, \dots, i_N, j_2, \dots, j_N, a_t\} & r_t = 0 \end{cases}$ 
10:    Find competitor item  $a_t^C$  of  $a_t$ 
11:    Store transition  $(s_t, a_t, r_t, s_{t+1}, a_t^C / \text{null})$  in  $\mathcal{D}$ 
12:    Sample minibatch of transitions  $(s, a, r, s', a^C / \text{null})$  from  $\mathcal{D}$ 
13:    Set  $y = \begin{cases} r & \text{terminal } s' \\ r + \gamma \max_{a'} Q(s', a'; \theta) & \text{non-terminal } s' \end{cases}$ 
14:    if  $a^C$  exists then
15:      Minimize  $(y - Q(s, a; \theta))^2 - \alpha (Q(s, a; \theta) - Q(s, a^C; \theta))^2$  according to Equ (11)
16:    else
17:      Minimize  $(y - Q(s, a; \theta))^2$  according to Equ (9)
18:    end if
19:  end for
20: end for

```

[5]; then the agent observes the reward r_t from users (line 8) and updates the state to s_{t+1} (line 9) and try to find the competitor item a_t^C (line 10); and finally the recommender agent stores transitions $(s_t, a_t, r_t, s_{t+1}, a_t^C / \text{null})$ into replay memory $\mathcal{D}$ (line 11). For model training stage: the recommender agent samples minibatch of transitions $(s, a, r, s', a^C / \text{null})$ from replay memory $\mathcal{D}$ (line 12), and then updates the parameters according to Equation (9) or Equation (11) (lines 13-18).

In the algorithm, we introduce widely used techniques to train our framework. For example, we utilize a technique known as *experience replay* [12] (lines 1,11,12), and introduce separated evaluation and target networks [16], which can help smooth the learning and avoid the divergence of parameters. Moreover, we leverage *prioritized sampling strategy* [18] to assist the framework learning from the most important historical transitions.

3.4 Offline Test

In the previous subsection, we have discussed how to train a DQN-based recommender model. Now we formally present the offline test of our proposed framework DEERS.

The intuition of our offline test method is that, for a given recommendation session, the recommender agent reranks the items in this session according to the items’ Q-value calculated by the trained DQN, and recommends the item with maximal Q-value to the user. Then the recommender agent observes the reward from users and updates the state. The reason why recommender agent only reranks items in this session rather than items in the whole item space is that for the historical offline dataset, we only have

the ground truth rewards of the existing items in this session. The offline test algorithm in details is presented in Algorithm 2.

Algorithm 2 Offline Test of DEERS Framework.

Input: Initial state s_1 , items $l_1, \dots, l_T$ and corresponding rewards $r_1, \dots, r_T$ of a session $\mathcal{L}$.

Output: Recommendation list with new order

```

1: for  $t = 1, T$  do
2:   Observe state  $s_t = \{i_1, \dots, i_N, j_1, \dots, j_N\}$ 
3:   Calculate Q-values of items in  $\mathcal{L}$ 
4:   Recommend item  $l_{max}$  with maximal Q-value
5:   Observe reward  $r_t$  from users (historical logs)
6:   Set  $s_{t+1} = \begin{cases} \{i_2, \dots, i_N, l_{max}, j_1, \dots, j_N\} & r > 0 \\ \{i_1, \dots, i_N, j_2, \dots, j_N, l_{max}\} & r = 0 \end{cases}$ 
7:   Remove  $l_{max}$  from  $\mathcal{L}$ 
8: end for

```

In each iteration of a test recommendation session $\mathcal{L}$, given the state s_t (line 2), the recommender agent recommends an item l_{max} with maximal Q-value calculated by the trained DQN (lines 3-4), then observes the reward r_t from users (line 5) and updates the state to s_{t+1} (line 6), and finally it removes item l_{max} from $\mathcal{L}$ (line 7). Without the loss of generality, the RA recommends one item to the user each time, while it is straightforward to extend it by recommending multiple items.

3.5 Online Test

We also do online test on a simulated online environment. The simulated online environment is also trained on users' logs, but not on the same data for training the DEERS framework. The simulator has the similar architecture with DEERS, while the output layer is a softmax layer that predicts the immediate feedback according to current state s_t and a recommended item a_t . We test the simulator on users' logs (not the data for training the DEERS framework and simulator), and experimental results demonstrate that the simulated online environment has overall 90% precision for immediate feedback prediction task. This result suggests that the simulator can accurately simulate the real online environment and predict the online rewards, which enables us to test our model on it.

When test DEERS framework on the well trained simulator, we can feed current state s_t and a recommended item a_t into the simulator, and receive immediate reward $r(s_t, a_t)$ from it. The online test algorithm in details is presented in Algorithm 3. Note that we can control the length of the session (the value T) manually.

In each iteration of a test recommendation session $\mathcal{L}$, given the state s_t (line 5), the RA recommends an item a_t following policy π (line 6), which is obtained in model training stage. Then the RA feed s_t and a_t into the simulated online environment and observes the reward r_t (line 7) and updates the state to s_{t+1} (line 8).

4 EXPERIMENTS

In this section, we conduct extensive experiments with a dataset from a real e-commerce site to evaluate the effectiveness of the proposed framework. We mainly focus on two questions: (1) how the

Algorithm 3 Online Test of DEERS Framework.

```

1: Initialize action-value function  $Q$  with well trained weights
2: for session = 1,  $M$  do
3:   Initialize state  $s_1$  from previous sessions
4:   for  $t = 1, T$  do
5:     Observe state  $s_t = \{i_1, \dots, i_N, j_1, \dots, j_N\}$ 
6:     Execute action  $a_t$  following policy  $\pi$ 
7:     Observe reward  $r_t$  from users (online simulator)
8:     Set  $s_{t+1} = \begin{cases} \{i_2, \dots, i_N, a_t, j_1, \dots, j_N\} & r_t > 0 \\ \{i_1, \dots, i_N, j_2, \dots, j_N, a_t\} & r_t = 0 \end{cases}$ 
9:   end for
10: end for

```

proposed framework performs compared to representative baselines; and (2) how the negative feedback (skipped items) contribute to the performance. We first introduce experimental settings. Then we seek answers to the above two questions. Finally, we study the impact of important parameters on the performance of the proposed framework.

4.1 Experimental Settings

We evaluate our method on a dataset of September, 2017 from JD.com. We collect 1,000,000 recommendation sessions (9,136,976 items) in temporal order, and use first 70% as training set and other 30% as test set. For a given session, the initial state is collected from the previous sessions of the user. In this paper, we leverage $N = 10$ previously clicked/ordered items as positive state and $N = 10$ previously skipped items as negative state. The reward r of skipped/clicked/ordered items are empirically set as 0, 1, and 5, respectively. The dimension of the embedding of items is 50, and we set the discounted factor $\gamma = 0.95$. For the parameters of the proposed framework such as α and γ , we select them via cross-validation. Correspondingly, we also do parameter-tuning for baselines for a fair comparison. We will discuss more details about parameter selection for the proposed framework in the following subsections.

For the architecture of Deep Q-network, we leverage a 5-layer neural network, in which the first 3 layers are separated for positive and negative signals, and the last 2 layers connects both positive and negative signals, and outputs the Q-value of a given state and action.

As we consider our offline test task as a reranking task, we select MAP [29] and NDCG@40 [9] as the metrics to evaluate the performance. The difference of ours from traditional Learn-to-Rank methods is that we rerank both clicked and ordered items together, and set them by different rewards, rather than only rerank clicked items as that in Learn-to-Rank problems. For online test, we leverage the accumulated rewards in the session as the metric.

4.2 Performance Comparison for Offline Test

We compare the proposed framework with the following representative baseline methods:

- **CF:** Collaborative filtering[3] is a method of making automatic predictions about the interests of a user by collecting preference information from many users, which is based on

Figure 5: Overall performance comparison in offline test.

the hypothesis that people often get the best recommendations from someone with similar tastes to themselves.

- **FM:** Factorization Machines[22] combine the advantages of support vector machines with factorization models. Compared with matrix factorization, higher order interactions can be modeled using the dimensionality parameter.
- **GRU:** This baseline utilizes the Gated Recurrent Units (GRU) to predict what user will click/order next based on the clicking/ordering histories. To make comparison fair, it also keeps previous $N = 10$ clicked/ordered items as states.
- **DEERS-p:** We use a Deep Q-network[16] with embeddings of users' historical clicked/ordered items (state) and a recommended item (action) as input, and train this baseline following Eq.(2). Note that the state is also captured by GRU.

The results are shown in Figure 5. We make following observations:

- Figure 5 (a) illustrates the training process of DEERS. We can observe that the framework approaches to convergence when the model are trained by 500,000 sessions. The fluctuation of the curve causes by the parameter replacement from evaluation DQN to target DQN.
- CF and FM achieve worse performance than GRU, DQN and DEERS, since CF and FM ignore the temporal sequence of the users' browsing history, while GRU can capture the temporal sequence, and DEERS-p and DEERS are able to continuously update their strategies during the interactions.
- GRU achieves worse performance than DEERS-p, since we design GRU to maximize the immediate reward for recommendations, while DEERS-p is designed to achieve the trade-off between short-term and long-term rewards. This result suggests that introducing reinforcement learning can improve the performance of recommendations.
- DEERS performs better than DEERS-p because we integrate both positive and negative items (or feedback) into DEERS, while DEERS-p is trained only based on positive items. This result indicates that negative feedback can also influence

Figure 6: Overall performance comparison in online test.

the decision making process of users, and integrating them into the model can lead to accurate recommendations.

To sum up, we can draw answers to the two questions: (1) the proposed framework outperforms representative baselines; and (2) negative feedback can help the recommendation performance.

4.3 Performance Comparison for Online Test

We do online test on the aforementioned simulated online environment, and compare with DEERS with GRU and several variants of DEERS. Note that we do not include CF and FM baselines as offline test since CF and FM are not applicable to the simulated online environment.

- **GRU:** The same GRU as in the above subsection.
- **DEERS-p:** The same DEERS-p as in the above subsection.
- **DEERS-f:** This variant is a traditional 5-layer DQN where all layers are fully connected. Note that the state is also captured by GRU.
- **DEERS-t:** In this variant, we remove GRU units and just concatenate the previous $N = 10$ clicked/ordered items as positive state and previous $N = 10$ skipped items as negative state.
- **DEERS-r:** This variant is to evaluate the performance of pairwise regularization term, so we set $\alpha = 0$ to eliminate the pairwise regularization term.

As the test stage is based on the simulator, we can artificially control the length of recommendation sessions to study the performance in short and long sessions. We define short sessions with 100 recommended items, while long sessions with 300 recommended items. The results are shown in Figure 6. It can be observed:

- DEERS outperforms DEERS-f, which demonstrates that the new architecture of DQN can indeed improve the recommendation performance.
- An interesting discovery is that DEERS-f achieves even worse performance than DEERS-p, which is trained only based on positive feedback. This result indicates that a full connected DQN cannot detect the differences between positive and negative items, and simply concatenating positive and negative feedback as input will reduce the performance of traditional DQN. Thus redesigning DQN architecture is a necessary.
- In short recommendation sessions, GRU and DEERS-p achieve comparable performance. In other words, GRU and reinforcement learning models like DEERS-p can both recommend proper items matching users' short-term interests.

Figure 7: Parameter sensitiveness. (a) α that controls the pairwise regularization term. (b) N that controls the length of state.

- In long recommendation sessions, DEERS-p outperforms GRU significantly, because GRU is designed to maximize the immediate reward for recommendations, while reinforcement learning models like DEERS-p are designed to achieve the trade-off between short-term and long-term rewards.
- DEERS outperforms DEERS-t and DEERS-r, which suggests that introducing GRU to capture users' sequential preference and introducing pairwise regularization term can improve the performance of recommendations.

In summary, appropriately redesigning DQN architecture to incorporate negative feedback, leveraging GRU to capture users' dynamic preference and introducing pairwise regularization term can boost the recommendation performance.

4.4 Parameter Sensitivity Analysis

Our method has two key parameters, i.e., α that controls the pairwise regularization term and N that controls the length of state. To study the impact of these parameters, we investigate how the proposed framework DEERS works with the changes of one parameter, while fixing other parameters.

Figure 7 (a) shows the parameter sensitivity of α in offline recommendation task. The performance for recommendation achieves the peak when $\alpha = 0.1$. In other words, the pairwise regularization term indeed improves the performance of the model; however, it is less important than the first term $(y - Q(s_+, s_-, a; \theta))^2$ in Equation 10.

Figure 7 (b) demonstrates the parameter sensitivity of N in offline recommendation task. We find that with the increase of N , the performance improves, and the performance is more sensitive with positive items. In other words, user's decision mainly depends on the items she/he clicked or ordered, but the skipped items also influence the decision making process.

5 RELATED WORK

In this section, we briefly review works related to our study. In general, the related work can be mainly grouped into the following categories.

The first category related to this paper is traditional recommendation techniques. Recommender systems assist users by supplying a list of items that might interest users. Efforts have been made on offering meaningful recommendations to users. Collaborative filtering [13] is the most successful and the most widely used technique, which is based on the hypothesis that people often get the best recommendations from someone with similar tastes to themselves [3]. Another common approach is content-based filtering [17], which tries to recommend items with similar properties to those that a user ordered in the past. Knowledge-based systems [1] recommend items based on specific domain knowledge about how certain item features meet users needs and preferences and how the item is useful for the user. Hybrid recommender systems are based on the combination of the above mentioned two or more types of techniques [4]. The other topic closely related to this category is deep learning based recommender system, which is able to effectively capture the non-linear and non-trivial user-item relationships, and enables the codification of more complex abstractions as data representations in the higher layers [32]. For instance, Nguyen et al. [20] proposed a personalized tag recommender system based on CNN. It utilizes constitutional and max-pooling layer to get visual features from patches of images. Wu et al. [31] designed a session-based recommendation model for real-world e-commerce website. It utilizes the basic RNN to predict what user will buy next based on the click histories. This method helps balance the trade off between computation costs and prediction accuracy.

The second category is about reinforcement learning for recommendations, which is different with the traditional item recommendations. In this paper, we consider the recommending procedure as sequential interactions between users and recommender agent; and leverage reinforcement learning to automatically learn the optimal recommendation strategies. Indeed, reinforcement learning have been widely examined in recommendation field. The MDP-Based CF model can be viewed as approximating a partial observable MDP (POMDP) by using a finite rather than unbounded window of past history to define the current state [25]. To reduce the high computational and representational complexity of POMDP, three strategies have been developed: value function approximation [7], policy based optimization [19, 21], and stochastic sampling [10]. Furthermore, Mahmood et al. [15] adopted the reinforcement learning technique to observe the responses of users in a conversational recommender, with the aim to maximize a numerical cumulative reward function modeling the benefit that users get from each recommendation session. Taghipour et al. [27, 28] modeled web page recommendation as a Q-Learning problem and learned to make recommendations from web usage data as the actions rather than discovering explicit patterns from the data. The system inherits the intrinsic characteristic of reinforcement learning which is in a constant learning process. Sunehag et al. [26] introduced agents that successfully address sequential decision problems with high-dimensional combinatorial state and action spaces.

Zhao et al. [33, 34] propose a novel page-wise recommendation framework based on deep reinforcement learning, which can optimize a page of items with proper display based on real-time feedback from users.

6 CONCLUSION

In this paper, we propose a novel framework DEERS, which models the recommendation session as a Markov Decision Process and leverages Reinforcement Learning to automatically learn the optimal recommendation strategies. Reinforcement learning based recommender systems have two advantages: (1) they can continuously update strategies during the interactions, and (2) they are able to learn a strategy that maximizes the long-term cumulative reward from users. Different from previous work, we leverage Deep Q-Network and integrate skipped items (negative feedback) into reinforcement learning based recommendation strategies. Note that we design a novel architecture of DQN to capture both positive and negative feedback simultaneously. We evaluate our framework with extensive experiments based on data from a real e-commerce site. The results show that (1) our framework can significantly improve the recommendation performance; and (2) skipped items (negative feedback) can assist item recommendation.

There are several interesting research directions. First, in addition to positional order of items we used in this work, we would like to investigate more orders like temporal order. Second, we would like to validate with more agent-user interaction patterns, e.g., storing items into shopping cart, and investigate how to model them mathematically for recommendations. Finally, the items skipped by users may not be caused by users disliking them, but just not preferring as more as the items clicked/ordered or not viewing them in details at all. The week/wrong negative feedback may not improve or even reduce the performance when we consider the negative feedback. To capture stronger negative feedback, more information like dwell time can be recorded in users' behavior log and used in our framework.

ACKNOWLEDGEMENTS

This material is based upon work supported by, or in part by, the National Science Foundation (NSF) under grant number IIS-1714741 and IIS-1715940.

REFERENCES

- [1] Rajendra Akerkar and Priti Sajja. 2010. *Knowledge-based systems*. Jones & Bartlett Publishers.
- [2] Richard Bellman. 2013. *Dynamic programming*. Courier Corporation.
- [3] John S Breese, David Heckerman, and Carl Kadie. 1998. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 43–52.
- [4] Robin Burke. 2002. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction* 12, 4 (2002), 331–370.
- [5] Thomas Degris, Martha White, and Richard S Sutton. 2012. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839* (2012).
- [6] Georges E Dupret and Benjamin Piwowarski. 2008. A user browsing model to predict search engine click data from past observations. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 331–338.
- [7] Milos Hauskrecht. 1997. Incremental methods for computing bounds in partially observable Markov decision processes. In *AAAI/IAAI*. 734–739.
- [8] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [9] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [10] Michael Kearns, Yishay Mansour, and Andrew Y Ng. 2002. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine learning* 49, 2 (2002), 193–208.
- [11] Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems*. 2177–2185.
- [12] Long-Ji Lin. 1993. *Reinforcement learning for robots using neural networks*. Technical Report. Carnegie-Mellon Univ Pittsburgh PA School of Computer Science.
- [13] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing* 7, 1 (2003), 76–80.
- [14] Tariq Mahmood and Francesco Ricci. 2007. Learning and adaptivity in interactive recommender systems. In *Proceedings of the ninth international conference on Electronic commerce*. ACM, 75–84.
- [15] Tariq Mahmood and Francesco Ricci. 2009. Improving recommender systems with adaptive conversational strategies. In *Proceedings of the 20th ACM conference on Hypertext and hypermedia*. ACM, 73–82.
- [16] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [17] Raymond J Mooney and Loriene Roy. 2000. Content-based book recommending using learning for text categorization. In *Proceedings of the fifth ACM conference on Digital libraries*. ACM, 195–204.
- [18] Andrew W Moore and Christopher G Atkeson. 1993. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning* 13, 1 (1993), 103–130.
- [19] Andrew Y Ng and Michael Jordan. 2000. PEGASUS: A policy search method for large MDPs and POMDPs. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 406–415.
- [20] Hanh TH Nguyen, Martin Wistuba, Josif Grabocka, Lucas Rego Drumond, and Lars Schmidt-Thieme. 2017. Personalized Deep Learning for Tag Recommendation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 186–197.
- [21] Pascal Poupart and Craig Boutilier. 2005. VDCBPI: an approximate scalable algorithm for large POMDPs. In *Advances in Neural Information Processing Systems*. 1081–1088.
- [22] Steffen Rendle. 2010. Factorization machines. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*. IEEE, 995–1000.
- [23] Paul Resnick and Hal R Varian. 1997. Recommender systems. *Commun. ACM* 40, 3 (1997), 56–58.
- [24] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. Springer, 1–35.
- [25] Guy Shani, David Heckerman, and Ronen I Brafman. 2005. An MDP-based recommender system. *Journal of Machine Learning Research* 6, Sep (2005), 1265–1295.
- [26] Peter Sunehag, Richard Evans, Gabriel Dulac-Arnold, Yori Zwols, Daniel Visentin, and Ben Coppin. 2015. Deep Reinforcement Learning with Attention for Slate Markov Decision Processes with High-Dimensional States and Actions. *arXiv preprint arXiv:1512.01124* (2015).
- [27] Nima Taghipour and Ahmad Kardan. 2008. A hybrid web recommender system based on q-learning. In *Proceedings of the 2008 ACM symposium on Applied computing*. ACM, 1164–1168.
- [28] Nima Taghipour, Ahmad Kardan, and Saeed Shiry Ghidary. 2007. Usage-based web recommendations: a reinforcement learning approach. In *Proceedings of the 2007 ACM conference on Recommender systems*. ACM, 113–120.
- [29] Andrew Turpin and Falk Scholer. 2006. User performance versus precision measures for simple search tasks. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 11–18.
- [30] Qingyun Wu, Hongning Wang, Liangjie Hong, and Yue Shi. 2017. Returning is Believing: Optimizing Long-term User Engagement in Recommender Systems. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 1927–1936.
- [31] Sai Wu, Weichao Ren, Chengchao Yu, Gang Chen, Dongxiang Zhang, and Jingbo Zhu. 2016. Personal recommendation using deep recurrent neural networks in NetEase. In *Data Engineering (ICDE), 2016 IEEE 32nd International Conference on*. IEEE, 1218–1229.
- [32] Shuai Zhang, Lina Yao, and Aixin Sun. 2017. Deep Learning based Recommender System: A Survey and New Perspectives. *arXiv preprint arXiv:1707.07435* (2017).
- [33] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep Reinforcement Learning for Page-wise Recommendations. *arXiv preprint arXiv:1805.02343* (2018).
- [34] Xiangyu Zhao, Liang Zhang, Zhuoye Ding, Dawei Yin, Yihong Zhao, and Jiliang Tang. 2017. Deep Reinforcement Learning for List-wise Recommendations. *arXiv preprint arXiv:1801.00209* (2017).