



Adversarial Robustness of Deep Reinforcement Learning Based Dynamic Recommender Systems

Siyu Wang^{1†}, Yuanjiang Cao^{1†}, Xiaocong Chen¹, Lina Yao^{1*}, Xianzhi Wang² and Quan Z. Sheng³

¹ School of Computer Science and Engineering, University of New South Wales, Sydney, NSW, Australia, ² School of Computer Science, University of Technology Sydney, Sydney, NSW, Australia, ³ Department of Computing, Macquarie University, Sydney, NSW, Australia

OPEN ACCESS

Edited by:

Alejandro Bellogin,
Autonomous University of Madrid,
Spain

Reviewed by:

Yanghua Xiao,
Fudan University, China
Chaochao Chen,
Zhejiang University, China

*Correspondence:

Lina Yao
lina.yao@unsw.edu.au

[†]These authors have contributed equally to this work and share first authorship

Specialty section:

This article was submitted to
Recommender Systems,
a section of the journal
Frontiers in Big Data

Received: 26 November 2021

Accepted: 09 March 2022

Published: 03 May 2022

Citation:

Wang S, Cao Y, Chen X, Yao L,
Wang X and Sheng QZ (2022)
Adversarial Robustness of Deep
Reinforcement Learning Based
Dynamic Recommender Systems.
Front. Big Data 5:822783.
doi: 10.3389/fdata.2022.822783

Adversarial attacks, e.g., adversarial perturbations of the input and adversarial samples, pose significant challenges to machine learning and deep learning techniques, including interactive recommendation systems. The latent embedding space of those techniques makes adversarial attacks challenging to detect at an early stage. Recent advance in causality shows that counterfactual can also be considered one of the ways to generate the adversarial samples drawn from different distribution as the training samples. We propose to explore adversarial examples and attack agnostic detection on reinforcement learning (RL)-based interactive recommendation systems. We first craft different types of adversarial examples by adding perturbations to the input and intervening on the casual factors. Then, we augment recommendation systems by detecting potential attacks with a deep learning-based classifier based on the crafted data. Finally, we study the attack strength and frequency of adversarial examples and evaluate our model on standard datasets with multiple crafting methods. Our extensive experiments show that most adversarial attacks are effective, and both attack strength and attack frequency impact the attack performance. The strategically-timed attack achieves comparative attack performance with only 1/3 to 1/2 attack frequency. Besides, our white-box detector trained with one crafting method has the generalization ability over several other crafting methods.

Keywords: recommender systems (RS), deep reinforcement learning (deep RL), adversarial attack, robustness, deep learning—artificial neural network (DL-ANN)

1. INTRODUCTION

Recommendation systems are an effective means of alleviating information overload for Internet users. They generally filter out those less irrelevant ones from massive items of choice to improve user experience in multiple scenarios. Traditional recommendation systems extract features about user preferences, items, and users' past interactions with items to conduct content-based, collaborative, or hybrid recommendations (G.Adomavicius and A.Tuzhilin, 2005; Zhang et al., 2019). These models have not considered the changes in user preferences over time. In this regard, interactive recommendation systems emerge to capture personalized user preference dynamics. Generally, interactive recommendation systems cater to users' dynamic and personalized requirements by improving the rigid strategy of conversational recommendation

systems (Thompson et al., 2004; Mahmood and Ricci, 2007; Taghipour and Kardan, 2008). In recent years, they have been attracting increasing attention and employed in leading companies (e.g., Amazon, Netflix, and YouTube) for personalized recommendations.

Interactive recommendation systems can be considered a decision-making process where the system chooses an optimal action in each discrete step to maximize the user response evaluation. Common practices to model the interactions between recommendation systems and users include Multi-Armed Bandit (MAB) or Reinforcement Learning (RL). The former views the action choice as a repeated single process while the latter considers immediate and future rewards to better model long-term user preference behaviors. In RL-based recommender systems, a Markov Decision Process (MDP) agent estimates the value function by both action and state rather than merely by action as done by MAB.

However, small disturbances in the input data may fool the above practices (Szegedy et al., 2013; Goodfellow et al., 2015). Small imperceptible noises, such as adversarial examples, may increase prediction error or reduce reward in supervised and RL tasks—the input noise can be transferred to attack different parameters even different models, including recurrent network and RL (Huang et al., 2017; Gao et al., 2018). Besides, the vector representations of entity/relation embedding of the input of RL-based recommendation models make it challenging for humans to tell the true value or dig out the real issues in the models.

Recently, Browne and Swift (2020) point out that counterfactual reasoning can be used to generate adversarial samples. From the perspective of causal inference, one way to leverage counterfactual reasoning is by intervening on some causes in the data generation process to generate adversarial samples. Since both perturbations and counterfactual reasoning target the state space by introducing noise, attackers can easily leverage such characteristics of embedding vectors to disrupt recommendation systems silently. Therefore, it is crucial to study attack and defense methods for RL-based recommendation systems.

This study aims to develop a general detection model to detect attacks and increase the defense ability and robustness, which provides a practical strategy to overcome the dynamic “arm-race” of attack and defense in the long run. The problem is nontrivial due to three reasons. First, online attacks are inherently difficult to track or predict. Second, man-in-middle methods can attack the interactions between recommendation systems and users in web applications, giving opportunities for malicious people to disrupt recommendation systems in either a white-box or a black-box way. Third, the vast number of actions in RL-based recommendation systems poses a barrier to detecting user feedback since the exhaustively numerous items and user embedding vectors are not feasible to find the abnormal inputs. We propose an attack-agnostic detection model against adversarial examples for RL-based recommendation systems to overcome the above challenges. To the best of our knowledge, this is the study that focuses on the adversarial detection of RL-based Recommendation Systems. We make the following contributions:

- We systematically investigate different types of adversarial attacks and detection approaches focusing on RL-based recommendation systems and demonstrate the effectiveness of the designed adversarial examples and strategically-timed attacks.
- We propose an encoder-classification detection model for attack-agnostic detection. The encoder captures the temporal relationship among sequence actions in RL. We further use an attention-based classifier to highlight the critical time steps out of a large interactive space.
- We empirically show that even small perturbations or counterfactual states can significantly reduce most attack methods’ performance. Our statistical validation shows that multiple attack methods generate similar actions of the attacked system, providing insights for improving the efficacy of the detection performance.

2. RELATED STUDY

RL-based interactive recommendation. RL is a popular approach to the interactive recommendation. Traditional research applies Q-learning (Taghipour et al., 2007; Taghipour and Kardan, 2008) and MDP (Mahmood and Ricci, 2009) to web recommendation and conversational recommendation problems. Mahmood and Ricci (2007) first introduce reinforcement learning into an interactive recommendation by modifying MDP. Since then, deep learning has inspired more interest in the interactive recommendation. For example, Christakopoulou et al. (2018) employ reinforcement learning to improve feedback quality in interactive recommendation; Chen et al. (2019) adopt policy gradient to improve the scalability of interactive recommendation.

Adversarial attacks. We explore the test-time white-box attack for the RL-based recommender system. This branch of study starts from Szegedy et al. (2013), which first find that hardly perceptible perturbation can cause erroneous outputs of a convolutional neural network on image classification tasks. Goodfellow et al. (2015) exploit this topic further and incorporate the Fast Gradient Sign Method to attack neural networks; (Papernot et al., 2016) proposes Jacobian-based Saliency Map Attack (JSMA) algorithm to greedily select attack pixels by Jacobian matrix. Another view is to model the attack as an iterated optimization process, like the Deepfool model (Moosavi-Dezfooli et al., 2016) and PGD (Kurakin et al., 2016). Croce and Hein (2020) use the approximated hyperplane as Deepfool and guarantee the perturbed data point is closed to the hyperplane. Yu et al. (2021) exploits the feature level attack to achieve decent results. Chen and Gu (2020) factorize pixels into two variables and create a model that can control the sparsity of the attack. Specifically, Lin et al. (2017) design strategically-timed attacks and craft deceptive images to induce the agent to make the desired actions. Browne and Swift (2020) argue that counterfactual explanations produce adversarial examples in DL research, which modify the input to cause misclassification of the network. Huang et al. (2017) explore the adversarial attack deep Q network in video game playing and conclude that

retraining with adversarial examples can make the network more robust. Another thread of research applies adversarial attacks into environments for robust adversarial training. They either regard the attack as a destabilizing force to break the balance of agents in 3D scenarios (Pinto et al., 2017) or develop adversarial agents in multi-agent tasks during RL (Gleave et al., 2019). Generally, creating adversarial examples helps reduce the reward of on DQN and DDPG (Pattanaik et al., 2018), and a detection method can help better explore the potential of adversarial examples and make agents more robust in a dynamic process.

Adversarial example detection. Many adversarial detection methods are vulnerable to the loss functions targeted to fool them (Carlini and Wagner, 2017). Bendale et al. Bendale and Boulton (2016) present OpenMax to estimate the probability of data input from unknown classes explicitly. Since then, researchers have proposed a statistical approach (Hendrycks and Gimpel, 2016), binary classification approach (Metzen et al., 2017), outlier detection approach (Grosse et al., 2017), and history queries-based approach (Chen et al., 2020) to detect adversarial examples. Our study differs from Chen et al. (2020) in exploiting the nature of RL besides query-based white-box attacks. Detection models classify the benign samples and adversarial samples by discrepancy, which is verified in many areas (Cohen et al., 2020; Esmaeilpour et al., 2020; Vacanti and Van Looveren, 2020; Massoli et al., 2021), in this article, we exert the discrepancy in action space to detect adversarial samples, which turns out to be effective for multiple attack methods.

3. METHODOLOGY

This section introduces the components of an RL-based recommendation system, attack techniques that generate adversarial examples, and our scheme to detect white-box adversarial attacks. The overall structure can be found in **Figure 1**.

3.1. RL-Based Interactive Recommendation

Interactive recommendation systems suggest items to users and receive feedback. Given a user $u_j \in U = \{u_0, u_1, u_2, \dots, u_n\}$, a set of items $I = \{i_0, i_1, i_2, \dots, i_n\}$, and the user feedback history $i_{k_1}, i_{k_2}, \dots, i_{k_{t-1}}$, the recommendation system suggests a new item i_{k_t} . This problem can be regarded as an MDP, which comprises the following:

- State (s_t): a historical interaction between the user and the recommendation system computed by an embedding or encoder module.
- Action (a_t): an item or a set of items recommended by the RL agent.
- Reward (r_t): a variable related to user's feedback to guide the reinforcement model toward true user preference.
- Policy ($\pi(a_t|s_t)$): a probabilistic model consisting of an estimation and action generation parts. The training process aims to obtain an optimal policy for the recommendation.
- Value function ($Q(s_t, a_t)$): the agent's prediction of the reward of the current recommended item a_t .

The reinforcement agent could be an Actor–Critic Algorithm that consists of a critic network and an actor network (Xian et al., 2019). The attack model may generate adversarial examples using either the critic network (Huang et al., 2017) or the actor network (Pattanaik et al., 2018).

3.2. Attack Model

FGSM-based attack. We define an adversarial example as a little perturbation δ added onto the benign examples x to reduce the cumulative reward of an RL system. Suppose x is a sequence of feature vectors piped into RL model $\pi(s_t)$; x can be a composition of embedding vectors of users, relations, and items (Xian et al., 2019), or a feature vector encoded user and item information (Chen et al., 2019). Unlike perturbations on images or texts, δ can be large in interactive recommendation systems due to the enormous manual work to check the embedding vectors or feature vectors of massive users and items. We define an adversarial example as follows:

$$\begin{aligned} \min_{\delta} R_T &= \sum_{t=1}^T r_t, \\ r_t &= Q(s_t + \delta, a_t), \\ a_t &= \pi^*(a_t|s_t + \delta), \\ \text{subject to } &S(s_t, s_t + \delta) \leq l, 1 \end{aligned} \quad (1)$$

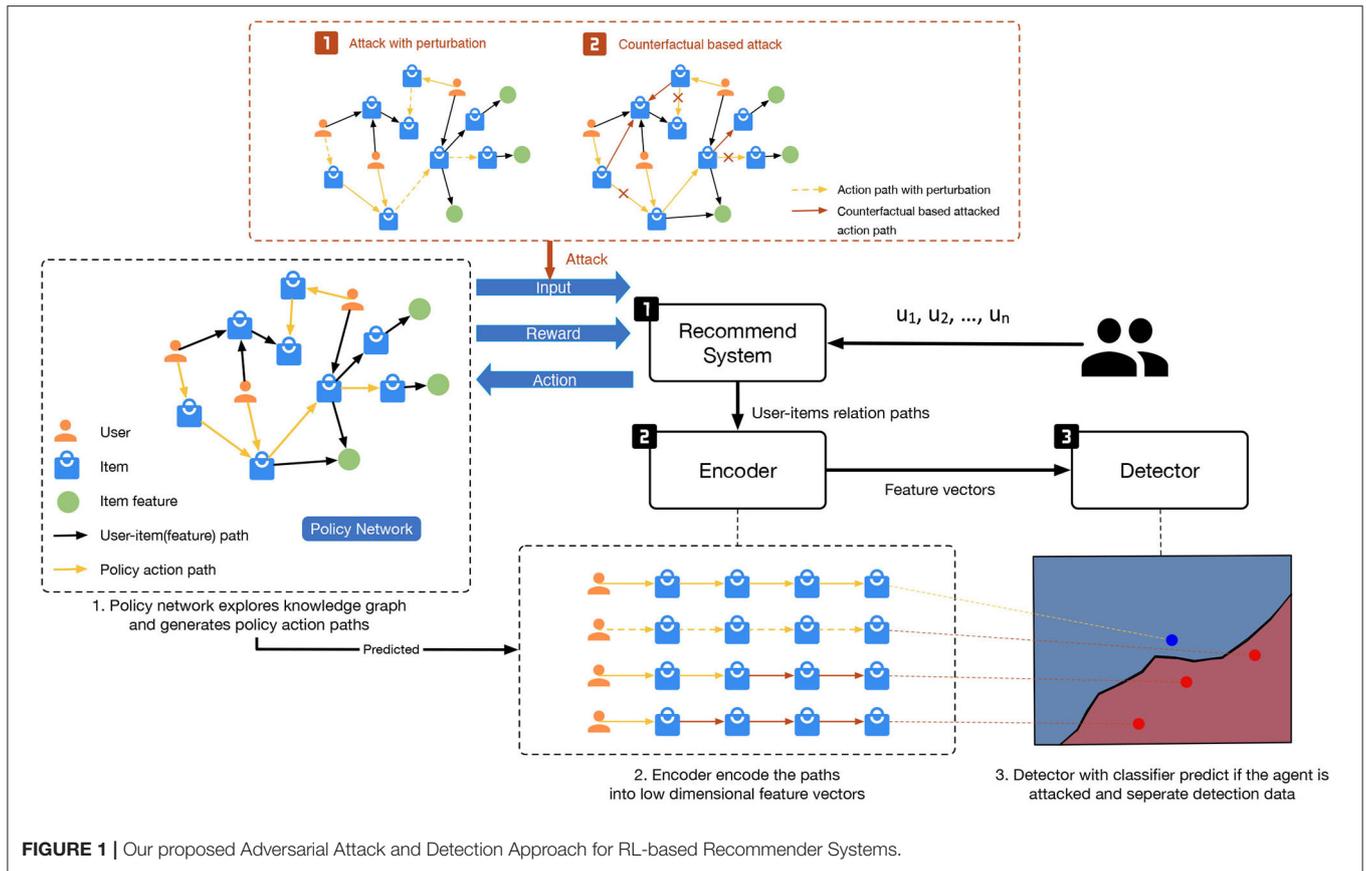
where R_T is the total reward of the recommendation agent, T is the length of a time step, π^* is the optimal policy learned by the training process, $S (< l)$ is a similarity metric that measures the distance between benign and adversarial examples. S is commonly defined as l_p bounded perturbation, or $|\delta|_p$ (Carlini et al., 2019). The computation of δ determines the method of attack. We aim to build a model with the generalization ability to detect examples from unknown adversarial distributions. Thus, we adopt three attack methods to validate the detection model performance: FGSM (Goodfellow et al., 2015) and its variant (Huang et al., 2017), JSMA (Papernot et al., 2016), and Deepfool (Moosavi-Dezfooli et al., 2016). FGSM can be presented as follows:

$$\delta_{inf} = \epsilon \text{sign}(\nabla_{s_t} J(Q_t, Q_{t-1} + r_t)), \quad (2)$$

where J is the loss function, Q_t is the critic function $Q(s_t, a_t)$. Optimizing J will lead to the critic value Q satisfying the Bellman equation. The FGSM method uses the gradient of the loss function, which can be computed efficiently, thus, requiring a small amount of additional computation.

To construct a detection model with the generalization ability, we train the detection model with FGSM examples and conduct the detection using other perturbation methods. We adopt the two norm variations in Huang et al. (2017) and define the norm constraint of perturbations as follows:

$$\begin{aligned} \delta_2 &= \epsilon \sqrt{d} * \frac{\nabla_{s_t} J(Q_t, Q_{t-1} + r_t)}{\|\nabla_{s_t} J(Q_t, Q_{t-1} + r_t)\|_2}, \\ \delta_1 &= \text{perturb highest } |\nabla_{s_t} J(Q_t, Q_{t-1} + r_t)| \text{dimension}. \end{aligned} \quad (3)$$



Attack with smaller frequency. The strategically-timed attack (Lin et al., 2017) aims to decrease the attack frequency without sacrificing the performance of the un-targeted reinforcement attack. We formally present it below:

$$\begin{aligned} \delta_t &= \delta_t * c_t \quad c_t \in \{0, 1\}, \\ \frac{\sum_{t=1}^T c_t}{T} &< d, \end{aligned} \tag{4}$$

where c_t is a binary variable that controls when to attack; $d < T$ is the frequency of adversarial examples. There are two approaches to generate the binary sequence $c_{1:T}$ optimizing a hard integer programming problem and generating sequences *via* heuristic methods. Let p_0, p_1 be the two maximum probability of a policy π , we define c_t as follows, which is different from Lin et al. (2017):

$$c_t = (p_0 - p_1) > threshold.$$

In our experiments, we let the RL-based recommendation system have a peak probability at the maximum action to test the importance of the action to attackers using the above formula. In contrast to the above methods, JSMA and Deepfool are based on the gradient of actions rather than the gradient of Q value. One key component of JSMA is saliency map computation used to decide which dimension of vectors (in Image classification is

pixels) are modified. Deepfool pinpoints the attack dimension by comparison of affine distances between some class and temporal classes. More details can be found in Papernot et al. (2016) and Moosavi-Dezfooli et al. (2016).

Counterfactual Based Attack. Counterfactual can find a similar version of the query input within some distributions, changing the model decisions and receiving a different classification. This helps to explain why a specific classification decision was made by a model and improve the interpretation of model boundaries (Yang et al., 2021), which is known as counterfactual explanations. Recent study reveals that counterfactual explanations produce adversarial examples in deep neural networks (Browne and Swift, 2020). Therefore, we propose to generate counterfactual user interacting processes to be the counterfactual-based attacks for the RL model.

Most of the adversarial examples are generated by adding perturbations. The counterfactual-based attack is recognized as one sub-type of adversarial example, which is different from traditional perturbations. One of the majority differences is that the counterfactual-based attack is generated by causal reasoning. With a casual relationship, we can perform interventions on causes to get counterfactual outcomes. To capture the casual relationships, we introduce Structural Causal Model(SCM) $M = \langle U, V, F \rangle$, given by a directed acyclic graph (DAG) \mathcal{G} , where:

- $U = (U_1, \dots, U_N)$ is a set of exogenous variables determined by unobserved and omitted factors. We assume that these noises are independent variables such that U_i is independent of all other noise variables.
- $V = (V_1, \dots, V_N)$ is a set of endogenous variables that are observed nodes in the DAG.
- $F = (f_1, \dots, f_N)$ is a set of structural equations representing the set of edges in the DAG. Each represents a causal mechanism that expresses the value of V_i as a function of the values of parents of V_i in \mathcal{G} and the noise term, that is $V_i = f_i(Pa(V_i)_i, U_i)$.

To simplify counterfactual reasoning, we assume that the input states follow the Local Causal Models (LCMs) (Pitis et al., 2020), stating that V_j is a parent node of V_i in \mathcal{G} if and only if there is a direct edge from V_j to V_i such that setting the value of V_j will have a direct effect on V_i through f_i . With this assumption, a large subspace \mathcal{L} often exists for each pair of nodes $(Pa(V_i)_j, V_i)$ in the DAG, in which two components are causally independent conditioning on a subset of parents nodes of V_i so that can be considered separately for training and inference.

Specifically, given two states with the same local factorization, we find the similar components of these two states. The similar components remain unchanged in the MDP process, representing that the critical components containing user identifiable information remain. Then, we test whether two sub-states without critical components, s_{ir} and s_{jr} , are locally independent. By performing interventions on one sub-state without critical components s_{jr} in SCM, we can calculate the new value of another sub-state s_{ir} . If the difference between the value after intervention and the original value of s_{ir} is within certain limits, setting the value of s_{jr} does not have a direct effect on s_{ir} through the causal mechanism. In that case, we conclude that the two sub-states without critical components are locally independent according to the LCMs. Leaving the critical components untouched, we produce a new counterfactual-based attack by swapping the two locally independent subsets of states s_i and s_j . The algorithm is given in Algorithm 1. This process can be interpreted by making intervention $do(S_t^{i,\dots,j}) = S_t^{j,\dots,i}$ on the LCMs $M^{\mathcal{L}}$ to obtain the simulation result (Pitis et al., 2020).

3.3. Detection Model

The detection model is a supervised classifier, which detects adversarial examples based on the actions of the reinforcement agent in a general feature space. Action-based detection exploits the fact that the defense can be constructed ignoring the attack type. Because state-based detection requires to model the distribution shift of various methods that increases the difficulty of modelling. Suppose the action distributions of an agent are shifted by adversarial examples (Section 4 shows statistical evidence of the drift). Given an abnormal action sequence $a = \pi^*(a|s + \delta)$ or a counterfactual action sequence, the detection model aims to establish a separating hyperplane between adversarial examples and normal examples, thereby measuring the probability $p(y|a, \theta)$ or $p(y|\pi^*, s, \delta, \theta)$, where y is a binary variable indicating whether the input data are attacked.

Algorithm 1 | Counterfactual State Generation.

```

1: procedure INDEPENDENT(SCM  $M$ , sub-state  $s_{ir}$ , sub-state  $s_{jr}$ )
2:   Set the value of  $s_{jr}$  to be  $s'_{jr}$   $\triangleright s'_{jr}$  has random value
3:   Calculate  $s'_{ir}$  by apply SCM  $M$  on  $s'_{jr}$ 
4:   return true if the  $s'_{ir}$  has the similar value with  $s_{ir}$   $\triangleright s_{ir}$  and  $s_{jr}$  are locally independent
5: end procedure

6: procedure GENERATION(SCM  $M$ , Current state  $s_i$ , a random state  $s_j$ )
7:   Find the similar component  $s_{same}$  in  $s_i$  and  $s_j$ 
8:    $s_{ir} \leftarrow s_i \setminus s_{same}$ 
9:    $s_{jr} \leftarrow s_j \setminus s_{same}$ 
10:  if INDEPENDENT( $M$ ,  $s_{ir}$ ,  $s_{jr}$ ) is true then
11:    Generate counterfactual state  $s'_i$  by swapping locally independent  $s_{ir}$  and  $s_{jr}$ 
12:  end if
13:  return counterfactual state  $s'_i$ 
14: end procedure

```

To detect the adversarial examples presented in the last section, we employ an attention-based classifier. We first conduct statistical analysis on the attacked actions whose result is shown in Section 4. The detection model consists of two parts. The first is an encoder to encode the action methods into a low-dimensional feature vector. The second is a classifier to separate different data. We adopt this encoder-decoder model to make a bottleneck and filter out noisy information. The formulation of GRU is as follows:

$$\begin{aligned}
 z_t &= \sigma_g(W_z a_t + U_z h_{t-1}), \\
 r_t &= \sigma_g(W_r a_t + U_r h_{t-1}), \\
 \hat{h}_t &= \tanh(W_h a_t + U_h \circ h_{t-1}), \\
 h_t &= (1 - z_t) \circ h_{t-1} + z_t \circ \hat{h}_t.
 \end{aligned} \tag{5}$$

We use an action sequence $a_1 : T$ to denote a series of user relation vectors or item embedding vectors and apply a recurrent model to encode the temporal relation into the feature vectors. We further adopt a single layer GRU network as our encoder and employ the attention-based dense net for detecting adversarial examples (formulated below).

$$\begin{aligned}
 \alpha_t &= \text{Softmax}(W_e e + b_e)_t, \\
 att, hid &= \sum_{t=1}^T \alpha_t h_t, \\
 p &= \text{Softmax}(W_{att} att + b_{att}),
 \end{aligned} \tag{6}$$

where e is the combined vector of action embedding and hidden states hid —we compute attention weights from embedding vectors and employ a linear unit to distribute probabilities to input time steps; h_t is the output of the encoder. The vectors processed through the attention layer are then piped into a linear unit with

softmax to compute the probability of adversarial examples. The loss function is the cross entropy between the true label and corresponding probability,

$$J(Att(a_{1:T}, y), y) = -y \circ \log(p).$$

4. EXPERIMENTS

In this section, we report our experiments to evaluate attack methods and our detection model. We first introduce the datasets and then provide quantitative evaluation and discussion on different attacks and our detection model.

4.1. Dataset and Experiment Setup

We conduct experiments based on two RL interactive recommendation systems. Following Chen et al. (2019) and Xian et al. (2019) over the real-world dataset – Amazon dataset (He and McAuley, 2016). This public dataset contains user reviews and metadata of the Amazon e-commerce platform from 1996 to 2014. We utilize three subsets named Beauty, Cellphones, and Clothing as our dataset. We directly use the dataset provided by Xian et al. (2019) on Github to reproduce their experiments. Details about Amazon dataset analysis can be found in Xian et al. (2019).

We conduct our attack and detection experiments based on (Xian et al., 2019). We preprocess the dataset by filtering out feature words with higher TF-IDF scores than 0.1. Then, we use 70% data in each dataset as the training set (and the rest as the test set) and the actions of the reinforcement agent as the detection data. We define the actions of PGPR (Xian et al., 2019) as heterogeneous graph paths that start from users and have a length of 4. The three Amazon sub-dataset (Beauty, Cellphones, and Clothing) contain 22,363, 27,879, and 39,387 users. To accelerate experiments, the first 10,000 users of each dataset are extracted for adversarial example production. Users in Beauty get on average 127.51 paths. The counterparts for Cellphones and Clothing are 121.92 and 122.71. We adopt the action file of l_∞ attack with an epsilon of 0.5 as the training set. As the number of paths is large, we utilize the first 1,00,000 paths for train and validation. The ratio of train validation is 80/20. Regarding the test, 1,00,000 paths from each action file are randomly sampled as the test set.

We slightly modify JSMA and Deepfool for our experiments—we create the saliency map by calculating the product of the target label and temporal label to achieve both effectiveness and higher efficiency (by 0.32 s per iteration) of JSMA; We also use sampling to decrease the computation load on a group of gradients for Deepfool. Besides, we set the hidden size of the GRU to 32 for the encoder, the drop rate of the attention-based classifier to 0.5, the maximum length of a user-item path to 4 [according to Xian et al. (2019)], and the learning rate and weight decay of the optimization solver, Adam, to $5e-4$ and 0.01, respectively.

4.2. Attack Experiments

This section reports our experiments on adversarial attacks. The first part shows the attack experiment results, followed by an analysis of the impact of attack frequency, attack intensity,

TABLE 1 | Adversarial attack results on Amazon Beauty, Cellphone, and Clothing.

Data	Method	Parameters	NDCG	Recall	HR	Precision
Beauty	Original	-	5.449	8.324	14.401	1.707
	FGSM l_1	$\epsilon = 0.1$	2.695	3.714	6.599	0.693
	FGSM l_2	$\epsilon = 1.0$	4.567	6.555	13.751	1.653
	FGSM l_{inf}	$\epsilon = 0.5$	2.830	3.909	7.351	0.787
	Counterfactual	-	5.324	8.089	14.077	1.658
	JSMA	-	2.984	3.844	8.254	0.931
Cellphone	Deepfool	-	3.280	4.352	9.548	1.050
	Original	-	3.559	5.510	9.031	0.933
	FGSM l_1	$\epsilon = 0.1$	2.146	3.007	4.832	0.493
	FGSM l_2	$\epsilon = 1.0$	3.000	4.444	7.471	0.784
Cloth	FGSM l_{inf}	$\epsilon = 0.5$	2.495	3.312	5.812	0.599
	Original	-	3.104	4.962	8.292	0.853
	FGSM l_1	$\epsilon = 0.1$	1.765	2.266	3.484	0.350
	FGSM l_2	$\epsilon = 1.0$	3.136	5.019	8.282	0.847
	FGSM l_{inf}	$\epsilon = 0.5$	1.449	2.005	2.865	0.286

TABLE 2 | MMD between a benign distribution and adversarial distribution on Amazon Beauty.

Data	Parameters	MMD-org	MMD- l_1
Original	-	0.121	0.620
FGSM l_1	$\epsilon = 0.1$	0.604	0.010
FGSM l_2	$\epsilon = 1.0$	0.016	0.573
FGSM l_{inf}	$\epsilon = 0.5$	0.570	0.011
Counterfactual	-	0.232	0.273
JSMA	-	0.412	0.034
Deepfool	-	0.177	0.458

and the action space of the recommendation system on the attack performance.

Adversarial attack results. We are interested in how vulnerable the agent is to perturbation in semantic embedding space. We consider an attack to be effective if a small perturbation leads to a notable performance reduction. We experimentally compare the performance of different attack methods (described in Section 3) in **Table 1**.

We reuse the evaluation metrics of the original model, namely Normalized Discounted Cumulative Gain (NDCG), Recall, Hit Ratio (HR), and Precision for evaluation on the amazon dataset. **Table 1** shows the attack results share the same trend with the distribution discrepancy in **Table 2**. Most attack methods significantly reduce the performance of the reinforcement system. FGSM l_1 achieves the best performance. It reveals that attacks on a single dimension can change the neural network's action drastically. Compared with l_1 and l_{inf} methods, FGSM l_2 is less effective on three datasets, where the evaluation metrics are mostly the same in contrast to the case without an attack (The original baseline in **Table 2**). It is worth mentioning that counterfactual attack does not perform well as the others. One of the possible reasons is that the generated counterfactual

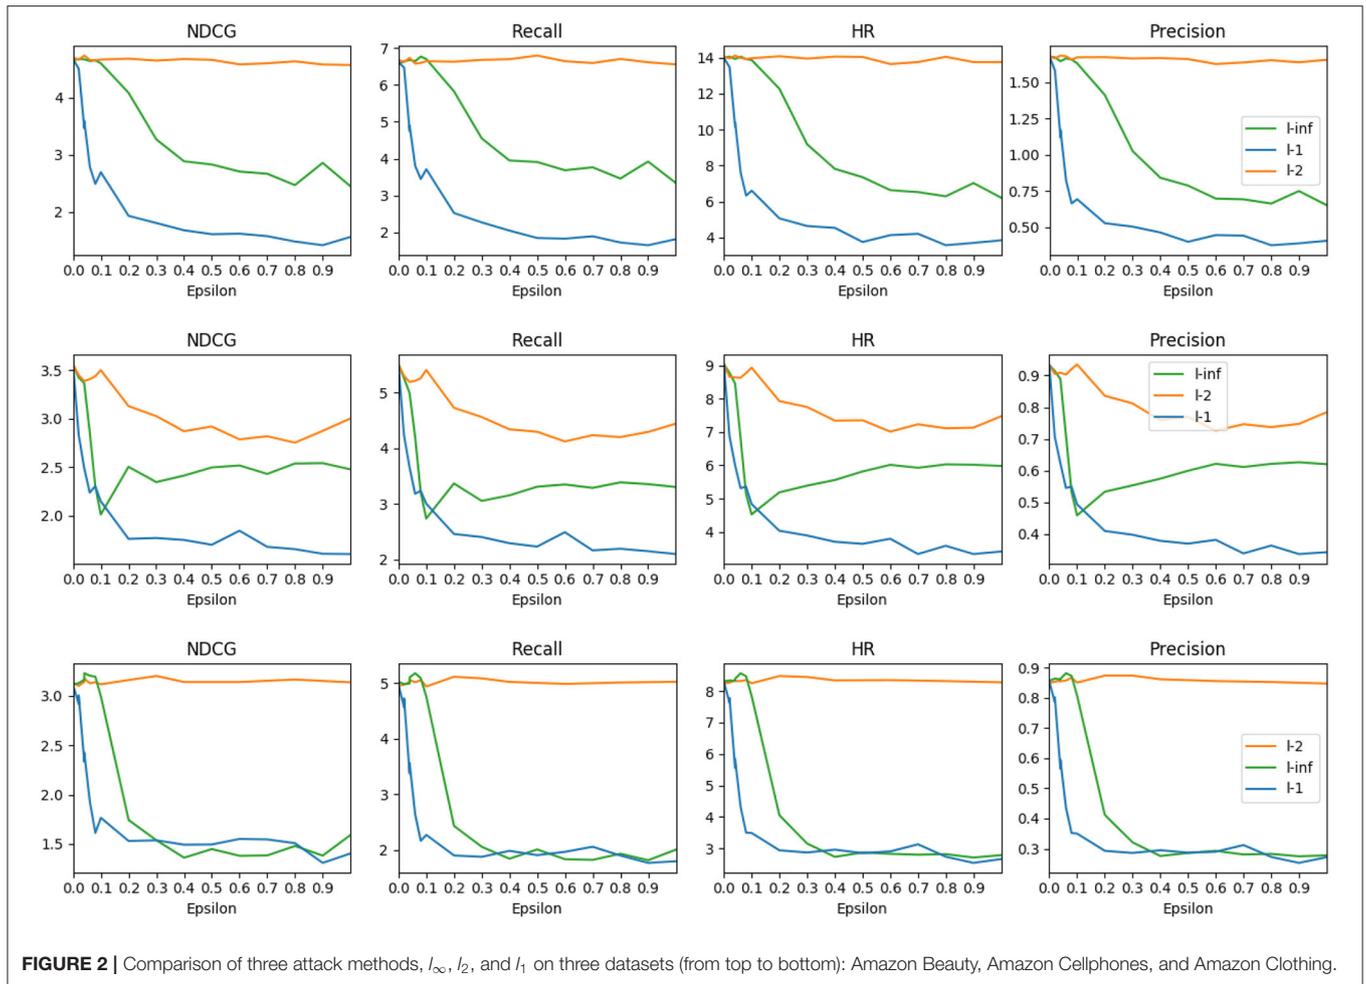


FIGURE 2 | Comparison of three attack methods, l_{∞} , l_2 , and l_1 on three datasets (from top to bottom): Amazon Beauty, Amazon Cellphones, and Amazon Clothing.

state still falls in the original latent space. The counterfactual attack can introduce noise to the current state by introducing irrelevant information from future states.

Impact of attack intensity. Adversarial examples make small perturbations to achieve notable changes in recommendation performance. Although larger perturbations on user-item interaction embeddings are not easily perceptible by humans, decreasing attack intensity might degrade attack effectiveness. To demonstrate the impact of different attack intensities in the context of RL-based recommender systems, we conduct the empirical experiment by varying the attack intensity, which is reflected by the ϵ parameter shown in Equation 2 and Equation 3. Experiment results of attack with epsilon variation of FGSM attack methods on three Amazon datasets (Figure 2) show that compared to a 0.0 value epsilon, all metric values decline as Epsilon increases, and l_1 attack achieves the best result. l_1 follows a similar yet more abrupt trend than the l_{∞} attack, while the l_2 attack achieves the worst performance regardless of the epsilon value. Huang et al. (2017) propose to attack the RL applied to games such as Atari. Their experiments reveal that the l_2 attack achieves comparable performance as l_1 and l_{inf} attacks do. To exclude the possibility that the

l_2 might be more effective with larger epsilon values, we set ϵ to 20 to test, but the result is the same. This observation reveals that the attack in user-item-feature embedding space shows different characteristics from attacks in the pixel space.

Another interesting observation is that the metric values show different trends depending on the datasets—unlike on Beauty and Cellphones, the l_{∞} attack achieves comparable performance to l_1 on the Clothing dataset when the ϵ is larger than 0.3. The result on the Cellphones dataset shows that the effectiveness of the l_{inf} attack diminishes as the ϵ continues increasing beyond 0.1.

Impact of attack frequency. We conduct two experiments on attack frequency, random attack, and strategic attack. In the random attack method, the adversarial examples are crafted with a frequency parameter, p_{freq} . In the strategically-timed attack, the adversarial examples are generated by the method shown in Section 3.2. The NDCG metric is presented in Figure 3; other metrics have a similar trend. It can be seen from Figure 3 that the random attack performs worse than the strategically-timed attack. Generating strategically adversarial examples in one-third to half time steps achieves a significant reduction in all metrics.

4.3. Detection Experiments

Analysis of adversarial examples. We use Maximum Mean Discrepancy as statistical measures of high dimensional data distribution distance. This divergence is defined as:

$$MMD(k, X_{org}, X_{adv}) = \sup_{k \in K} \left(\frac{1}{n} \sum_{i=1}^n k(x_{org,i}) - \frac{1}{m} \sum_{i=1}^m k(x_{adv,i}) \right),$$

where k is the kernel function, i.e., a radial basis function, which measures the distance between the means of two distributions (Table 2 shows the results); X_{org}, X_{adv} are benign and adversarial examples. The examples here refer to the actions generated by attacked RL agents. The distribution discrepancy of input data pierces through the deterministic model and shifts the action distribution that would be exploited by the detector. The data is randomly sampled from generated action embeddings of interactive recommendation systems. Each MMD is computed by averaging 40 batches of 500 samples. The actions generated by the RL agent are paths of the users-relation-item graph. As mentioned in Section 4.1, each user gets over 100 paths, which determines the overlapping of original data and adversarial examples decrease with the time step of the path growing. We choose the embedding of the last step to represent the recommended items.

MMD-org shows the discrepancy between the original and adversarial datasets, where $MMD-l_1$ presents the discrepancy between different attack methods. The results (Table 2) show that the adversarial distribution is different from the original distribution. Also, the disturbed distributions are closed to each other regardless of the attack type. This insight clarifies that we can use a classifier to separate benign data and adversarial data, and it can detect several attacks simultaneously, which might be transferred to other reinforcement learning attack detection tasks.

Detection Performance. From a statistical perspective, the above analysis shows that one classifier can detect multiple types of attacks. We evaluate the detection

TABLE 3 | Detection result and factor analysis.

Dataset	Attack	Precision	Recall	F1 Score
Beauty	l_1 0.1	0.919	0.890	0.904
	l_2 1.0	0.605	0.119	0.199
	l_{inf} 0.5	0.918	0.871	0.894
	Counterfactual	0.900	0.895	0.898
	JSMA	0.910	0.793	0.848
Cellphones	Deepfool	0.915	0.840	0.876
	l_1 0.1	0.801	0.781	0.791
	l_2 1.0	0.754	0.593	0.664
	l_{inf} 0.5	0.795	0.752	0.773
Clothing	l_1 1.0	0.810	0.825	0.817
	l_1 0.1	0.911	0.866	0.888
	l_2 1.0	0.541	0.099	0.168
	l_{inf} 0.5	0.912	0.879	0.895
Dataset	Frequency	Precision	Recall	F1 Score
Beauty	l_1 0.02	0.823	0.362	0.503
	l_1 0.04	0.906	0.754	0.823
	l_1 0.06	0.915	0.841	0.876
	l_1 0.08	0.918	0.872	0.894
	l_1 0.3	0.922	0.927	0.924
Dataset	Frequency	Precision	Recall	F1 Score
Beauty	l_1 0.579	0.921	0.912	0.917
	l_1 0.451	0.921	0.908	0.914
	l_1 0.316	0.918	0.879	0.898
	l_1 0.279	0.914	0.829	0.869
	l_1 0.118	0.837	0.401	0.543

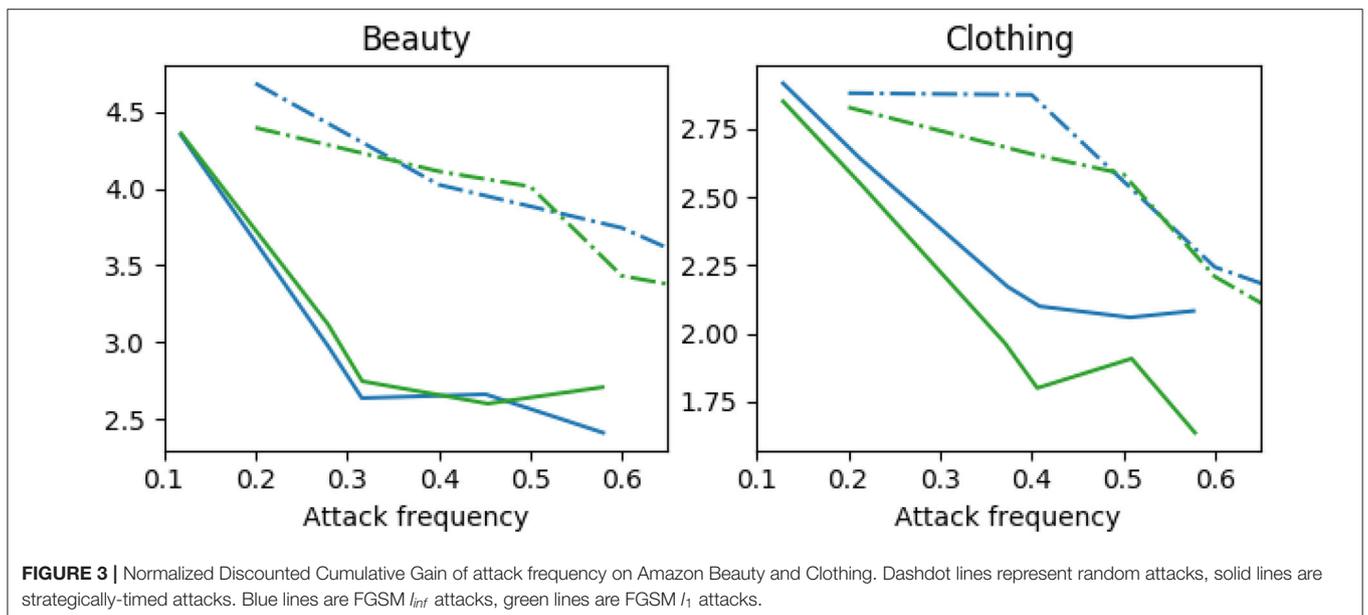


FIGURE 3 | Normalized Discounted Cumulative Gain of attack frequency on Amazon Beauty and Clothing. Dashdot lines represent random attacks, solid lines are strategically-timed attacks. Blue lines are FGSM l_{inf} attacks, green lines are FGSM l_1 attacks.

performance of different models using Precision, Recall, and F1 score.

We adopt an attention-based network for detection experiments. The detection model is trained on the FGSM l_1 attack with ϵ at 0.1 for all datasets. The results (Table 3) show that our detection model achieves better performance on attacks that cause serious disruption. The detection precision and recall rise as the attack is stronger. l_∞ attack validates this trend, which shows that our model can detect weaker attacks as well. The result of detection on l_2 attack can be reasoned with the MMD analysis shown above. High precision and low recall show that most l_2 adversarial examples are close to benign data, which confuses the detector. The l_1 attack with $\epsilon = 1.0$ validates our detector performs well yet achieves worse performance on other tests on the Cellphones dataset. Our model can also detect the counterfactual-based attack since the data distribution has been changed, verifying that our detection model can detect different types of attacks. Our results on factor analysis (Table 3) show that the detection model can detect attacks even under low attack frequencies. But the detection accuracy decreases as the attack frequency drops—the recall decreases significantly to 40.1% when 11.8% of examples represent attacks.

5. CONCLUSION

Adversarial attacks on reinforcement agents can greatly degrade user experience in interactive recommendation systems, as an intervention on causal factors can result in a different recommended result. In this article, we systematically

study adversarial attacks on RL-based recommendation systems by investigating different attack methods and the impact of attack intensity and frequency on the performance of adversarial examples. We conduct statistical analysis to show that classifiers, especially an attention-based detector, can well separate the detection data. Our extensive experiments show the excellent performance of our attack and detection models.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/Supplementary Materials, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

SW and YC write the original draft, conception, and design the experiments. XC provides the inspection and validation of the experiments. LY provides the supervision and manuscript reviewing. XW and QS provide manuscript reviewing. All authors contributed to the article and approved the submitted version.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fdata.2022.822783/full#supplementary-material>

REFERENCES

- Adomavicius G., and Tuzhilin, A. (2005). Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.* 17, 734–749. doi: 10.1109/TKDE.2005.99
- Bendale, A., and Boulton, T. E. (2016). “Towards open set deep networks,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Las Vegas, NV: IEEE), 1563–1572.
- Brown, K., and Swift, B. (2020). Semantics and explanation: why counterfactual explanations produce adversarial examples in deep neural networks. *arXiv preprint arXiv:2012.10076*.
- Carlini, N., Athalye, A., Papernot, N., Brendel, W., Rauber, J., Tsipras, D., et al. (2019). On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*.
- Carlini, N., and Wagner, D. (2017). “Adversarial examples are not easily detected: bypassing ten detection methods,” in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* (New York, NY: ACM), 3–14.
- Chen, H., Dai, X., Cai, H., Zhang, W., Wang, X., Tang, R., et al. (2019). “Large-scale interactive recommendation with tree-structured policy gradient,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33 (AAAI), 3312–3320. doi: 10.1609/aaai.v33i01.33013312
- Chen, J., and Gu, Q. (2020). “Rays: a ray searching method for hard-label adversarial attack,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (New York, NY), 1739–1747.
- Chen, S., Carlini, N., and Wagner, D. (2020). “Stateful detection of black-box adversarial attacks,” in *Proceedings of the 1st ACM Workshop on Security and Privacy on Artificial Intelligence*, 30–39.
- Christakopoulou, K., Beutel, A., Li, R., Jain, S., and Chi, E. H. (2018). “Q&R: a two-stage approach toward interactive recommendation,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (New York, NY: ACM), 139–148.
- Cohen, G., Sapiro, G., and Giryes, R. (2020). “Detecting adversarial samples using influence functions and nearest neighbors,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Seattle, WA), 14453–14462.
- Croce, F., and Hein, M. (2020). “Minimally distorted adversarial examples with a fast adaptive boundary attack,” in *International Conference on Machine Learning* (PMLR), 2196–2205.
- Esmailpour, M., Cardinal, P., and Koerich, A. L. (2020). “Detection of adversarial attacks and characterization of adversarial subspace,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Barcelona: IEEE), 3097–3101.
- Gao, J., Lanchantin, J., Soffa, M. L., and Qi, Y. (2018). “Black-box generation of adversarial text sequences to evade deep learning classifiers,” in *2018 IEEE Security and Privacy Workshops (SPW)* (San Francisco, CA: IEEE), 50–56.
- Gleave, A., Dennis, M., Wild, C., Kant, N., Levine, S., and Russell, S. (2019). “Adversarial policies: Attacking deep reinforcement learning,” in *International Conference on Learning Representations*.
- Goodfellow, I., Shlens, J., and Szegedy, C. (2015). “Explaining and harnessing adversarial examples,” in *International Conference on Learning Representations*. Available online at: <http://arxiv.org/abs/1412.6572>
- Grosse, K., Manoharan, P., Papernot, N., Backes, M., and McDaniel, P. (2017). On the (statistical) detection of adversarial examples. *arXiv[Preprint]*. arXiv: 1702.06280. Available online at: <https://arxiv.org/pdf/1702.06280.pdf>
- He, R., and McAuley, J. (2016). “Ups and downs: modeling the visual evolution of fashion trends with one-class collaborative filtering,” in *Proceedings of the 25th International Conference on World Wide Web* (Geneva: International World Wide Web Conferences Steering Committee), 507–517.

- Hendrycks, D., and Gimpel, K. (2016). Early methods for detecting adversarial images. *arXiv[Preprint]*. arXiv: 1608.00530. Available online at: <https://arxiv.org/pdf/1608.00530.pdf>
- Huang, S., Papernot, N., Goodfellow, I., Duan, Y., and Abbeel, P. (2017). Adversarial attacks on neural network policies. *arXiv[Preprint]*. arXiv: 1702.02284. Available online at: <https://arxiv.org/pdf/1702.02284.pdf>
- Kurakin, A., Goodfellow, I., and Bengio, S. (2016). Adversarial machine learning at scale. *arXiv preprint* arXiv:1611.01236.
- Lin, Y. -C., Hong, Z. -W., Liao, Y. -H., Shih, M. -L., Liu, M. -Y., and Sun, M. (2017). "Tactics of adversarial attack on deep reinforcement learning agents," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 3756–3762.
- Mahmood, T., and Ricci, F. (2007). "Learning and adaptivity in interactive recommender systems," in *Proceedings of the Ninth International Conference on Electronic Commerce* (New York, NY: ACM), 75–84.
- Mahmood, T., and Ricci, F. (2009). "Improving recommender systems with adaptive conversational strategies," in *Proceedings of the 20th ACM Conference on Hypertext and Hypermedia - HT '09* (New York, NY: ACM Press), 73.
- Massoli, F. V., Carrara, F., Amato, G., and Falchi, F. (2021). Detection of face recognition adversarial attacks. *Comput. Vis. Image Understand.* 202, 103103. doi: 10.1016/j.cviu.2020.103103
- Metzen, J. H., Genewein, T., Fischer, V., and Bischoff, B. (2017). On detecting adversarial perturbations. *arXiv[Preprint]*. arXiv: 1702.04267. Available online at: <https://arxiv.org/pdf/1702.04267.pdf>
- Moosavi-Dezfooli, S.-M., Fawzi, A., and Frossard, P. (2016). "Deepfool: a simple and accurate method to fool deep neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 2574–2582.
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. (2016). "The limitations of deep learning in adversarial settings," in *2016 IEEE European Symposium on Security and Privacy (EuroSec&P)* (Saarbruecken: IEEE), 372–387.
- Pattanaik, A., Tang, Z., Liu, S., Bommanna, G., and Chowdhary, G. (2018). "Robust deep reinforcement learning with adversarial attacks," in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems), 2040–2042.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. (2017). Robust adversarial reinforcement learning.
- Pitis, S., Creager, E., and Garg, A. (2020). Counterfactual data augmentation using locally factored dynamics. *arXiv preprint* arXiv:2007.02863.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013). Intriguing properties of neural networks.
- Taghipour, N., and Kardan, A. (2008). "A hybrid web recommender system based on q-learning," in *Proceedings of the 2008 ACM Symposium on Applied Computing* (New York, NY: ACM), 1164–1168.
- Taghipour, N., Kardan, A., and Ghidary, S. S. (2007). "Usage-based web recommendations: a reinforcement learning approach," in *Proceedings of the 2007 ACM Conference on Recommender Systems* (New York, NY: ACM), 113–120.
- Thompson, C. A., Goker, M. H., and Langley, P. (2004). A personalized system for conversational recommendations. *J. Artif. Intell. Res.* 21, 393–428. doi: 10.1613/jair.1318
- Vacanti, G., and Van Looveren, A. (2020). Adversarial detection and correction by matching prediction distributions. *arXiv preprint* arXiv:2002.09364.
- Xian, Y., Fu, Z., Muthukrishnan, S., de Melo, G., and Zhang, Y. (2019). Reinforcement knowledge graph reasoning for explainable recommendation. *arXiv preprint* arXiv:190.05237.
- Yang, F., Alva, S. S., Chen, J., and Hu, X. (2021). Model-based counterfactual synthesizer for interpretation. *arXiv preprint* arXiv:2106.08971.
- Yu, Y., Gao, X., and Xu, C.-Z. (2021). LAFEAT: piercing through adversarial defenses with latent features. *arXiv:2104.09284*.
- Zhang, S., Yao, L., Sun, A., and Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Comput. Surveys (CSUR)* 52, 5. 10.1145/3285029

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Wang, Cao, Chen, Yao, Wang and Sheng. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.