

Deep-attack over the deep reinforcement learning

Yang Li^{a,*}, Quan Pan^a, Erik Cambria^b

^aNorthwestern Polytechnical University, China

^bNanyang Technological University, Singapore

ARTICLE INFO

Article history:

Received 22 July 2021

Received in revised form 29 April 2022

Accepted 30 April 2022

Available online 10 May 2022

Keywords:

Adversarial attack

Deep reinforcement learning

Adversarial training

ABSTRACT

Recent adversarial attack developments have made reinforcement learning more vulnerable, and different approaches exist to deploy attacks against it, where the key is how to choose the right timing of the attack. Some work tries to design an attack evaluation function to select critical points that will be attacked if the value is greater than a certain threshold. This approach makes it difficult to find the right place to deploy an attack without considering the long-term impact. In addition, there is a lack of appropriate indicators of assessment during attacks. To make the attacks more intelligent as well as to remedy the existing problems, we propose the reinforcement learning-based attacking framework by considering the effectiveness and stealthy spontaneously, while we also propose a new metric to evaluate the performance of the attack model in these two aspects. Experimental results show the effectiveness of our proposed model and the goodness of our proposed evaluation metric. Furthermore, we validate the transferability of the model, and also its robustness under the adversarial training.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

With help of deep learning [1], deep reinforcement learning (DRL) has developed rapidly in the past few years and has become an important algorithm in many applications, e.g., automatic navigation [2], text generation [3], sentiment analysis [4], etc. In a deep reinforcement learning, it is to train an advanced policy network and guide the agent to take the optimal action in accordance with the state of the environment [5]. A policy network is usually a deep learning model to approximate the action-value function. However, the vulnerability of the deep learning model which is discovered by Szegedy et al. [6] has made DRL unstable when tackling different types of adversarial attacks. Attacks are usually deployed by adding some carefully crafted perturbations to the input and guiding the deep learning model to give incorrect outputs with high confidence. Unlike a typical supervised/unsupervised learning task which can be interpreted as a two-element tuple, i.e., $\langle \mathcal{X}, \mathcal{Y} \rangle$ (where \mathcal{X} indicates the data and \mathcal{Y} indicates the label), there are at least four elements i.e., $\langle S, A, \pi, R \rangle$ in a typical DRL (where A is the action set, S is the state set, R is the reward set, and π is the policy network). Therefore, DRL would be theoretically more vulnerable than supervised/unsupervised learning. And in fact, there have been many studies on attacking DRL by deploying perturbations on

different variables. In general, depending on the attack variables, they can be classified as reward-based attacks, observation-based attacks, action-based attacks, and policy-based attacks [7], and all of these attack strategies are designed to mislead the agent into taking the wrong action. Because it is not a one-time attack, an attack in DRL is more difficult than an attack in the general domain due to the series of steps involved in executing an action with a well-trained policy network. In addition, there is no clear definition of what constitutes a successful attack. Instead, there are two main constraints for a successful attack on an image, namely effectiveness and stealth. Therefore, it is necessary to propose evaluation metrics for the effectiveness of attacks on DRLs.

Attacking effectiveness means an attack confusing the model successfully with few training samples, however in DRL we need to consider both long-term success and short-term success as it is a series of actions. Specifically, short-term success means an attack affects the agent instantly and misleads it taking the wrong action immediately. This type of attack focuses on immediate effectiveness. While long-term success means an attack destroys the task with the fewest actions. This type of attack has been widely applied in many attacking algorithms [8,9]. The stealthiness of an attack is another important feature, which is important in practical applications because if an attack can be easily detected, it will simply be ended with brute force. A perfect attack is like the advanced persistent threat (APT) attack in network security, it will infiltrate the target for a long time until the final attack is carried out to destroy the mission.

* Corresponding author.

E-mail addresses: liyngnpu@nwpu.edu.cn (Y. Li), quanpan@nwpu.edu.cn (Q. Pan), liyngnpu@nwpu.edu.cn (E. Cambria).

However, there are seldom works combining the two attacks together. Generally, if we can find the key points for attacking which may have a long-term effect, the attack will be more effective. But the truth is that it is difficult to find them as they are not connected with future action directly. However, as we mentioned above, the short-term attack is an instant action, while the long-term attack is global action. If we can find the connection between these two types of attacks, their effectiveness will improve a lot. In this paper, to cope with such a problem, a general attacking framework (i.e., attacker) with Deep Q-Network (DQN) is proposed, the goal of the attacker policy network is to learn the timing of the attack. Generally, it is to find a policy π^* that maximizes the total rewards from all perturbed states, which is illustrated as follows:

$$\pi^* = \underset{r}{\operatorname{argmax}} \{E[\sum_{t=0}^{M_X} \gamma r^X(\pi(\operatorname{ATK}(s_t); \theta^X))]\} \quad (1)$$

where M_X is the total step and $r^X(\cdot)$ denotes the attacking rewards function, $\pi(\cdot)$ usually is a neural network which can be parameterized with θ^X , $\operatorname{ATK}(\cdot)$ is the attacking method, and γ is the discount factor. To ensure the stealthy and effectiveness of attacks, we designed the short-term reward r_{str} and long-term reward r_{ltr} to guide the agent, at the meantime, these two rewards can also be appropriate evaluation indicators. The contribution of this paper can be summarized as follows:

- Proposed a novel attacking framework which not only considers the long-term attack but also focuses on the short-term attack based on deep reinforcement learning;
- Proposed a new evaluation metric for the attacking effectiveness and stealth;
- Validated the effectiveness of the proposed model with empirical experiments.

This paper is organized as follows: Section 2 describes the related works about the adversarial attack in reinforcement learning. Section 3 provides the background of DQN and the traditional attacking method about critical point selection. Section 4 specifies how we design our attack method and Section 5 summarizes the evaluation results. Finally, we conclude our work and give future works in Section 6.

2. Related works

There are increasing researches on the adversarial attack and defense that exists in reinforcement learning, we will give a brief summarizing of existing works in terms of adversarial attack and adversarial defense separately.

2.1. Adversarial attack

There are plenty of works about the adversarial attack in DRL [10], and based on the attacking variables, they can be classified into the observation-based attack, reward-based attack, action-based attack, policy-based attack, and environment-based attack [7]. Observation-based attack means the attack occurs on the state observed by the agent, and the attack usually is deployed with the image-based attacking method. However, it is difficult to distinguish the difference between the observation- and environment-based attack as the state and environment can be the subset of the observation in general. For example, Huang et al. [11] who was the first one to explore the vulnerability in reinforcement learning deployed the attacks (i.e., Fast Gradient Sign Method (FGSM)) over the state directly. Behzadan et al. [12] deployed the attack over the state to induce the agent with a transfer learning-based method on the DQN. To find the right

time of attacking, Lin et al. [13] proposed the Strategically-Timed Attack which only selects the key framework to conduct the attack by designing a judgment function based on the state-action pair. The attacking method they deployed is Carlini and Wagner attack (C&W) [14]. Similarly, Kos et al. [15] adopted value function to guide the attack by judging whether the value is higher than a certain threshold, and Russo et al. [16] applied the Markov decision process to select the attacking frame to poison. All those attacks mentioned above need to know the parameters of the model, i.e., white-box attack. While there also are works on the black-box attack that can access only the input and output of the RL. For example, Huang et al. [11] utilized a surrogate policy network with a rollout technique to construct the adversarial state, and experimental results showed the usefulness of attack. Zhao et al. [17] proposed an approximation model to deploy the black-box attack by predicting the target agent's future actions based on their time-series information. And similar works can be found in works [12,18].

Apart from manipulating the environment or state that agents are in, some works try the decrease the performance by crafting the rewards, which is called reward-based attack listed in work [7]. For example, Tretschk et al. [8] adopted the adversarial transformer network to craft the adversarial rewards over the victim agent with a sequence of attacks, and the agent is misled in the opposite reward direction over time. While Han et al. [19] applied a more direct way to conduct the reward-based attack by flipping a certain number of rewards signs. Kiourti et al. [20] utilized the Trojan technique to attack the target policy network by deploying the in-band reward modification. Perturb over the action space is another approach to attack the agent, for example, Lee et al. [21] conducted the action space attack by reformulating the optimization problems that distribute the attacks across the action and temporal dimensions. In this paper, the goal is to learn an intelligent attacker which can deploy the attack at the right place by considering the effectiveness and stealthy spontaneously. Therefore, the environment is where we would like to deploy the attack.

2.2. Adversarial defense

All defense methods applied in reinforcement learning can be categorized into those three types, i.e., adversarial training, robust learning, and adversarial detection. Adversarial training is to add the adversarial samples into the training data, the main goal is to improve the generalization of the model. Since Goodfellow et al. [22] proposed FGSM to conduct the adversarial training over the image classification, this method has been successfully applied in reinforcement learning. For example, Kos et al. [15] and Pattanaik et al. [23] applied the FGSM to obtain adversarial examples constantly which helps the classification model acquire the resistance to adversarial examples. While the limitation of this method is that it cannot defend the outlier that does not exist in the training set. Apart from adding the adversarial samples constantly, some researchers do it differently, for example, Behzadan et al. [24] validated that if the adversarial samples generated by FGSM are added to the training set with a certain probability (i.e., P), better defense effects can also be achieved. And empirical results of DQN show that when setting P to 0.2 or 0.4, DQN can obtain a normal performance after the adversarial training. Based on this work, Behzadan et al. [25] adopted the ε -greedy search to find a more stable and efficient way to add adversarial samples in their most recent work. Another adversarial training approach is through the min-max game, For example, Gu et al. [26] formulated the task as a zero-sum dual-agent Markov game by splitting the agent into adversarial and protagonist, and the robustness is achieved from the min-max game between

them. Similarly, Pinto et al. [27] introduced two agents to play the zero-sum discount game to ensure the robustness of the policy learning. Different from the adversarial playing, Behzadan et al. [28] adopt an equivalent model as the noisy network to generate the adversarial samples via FGSM. Neklyudov et al. [29] applied the Gaussian variance layer to generate the adversarial samples, and empirical results show that this method is effective in improving the ability of exploration and robustness of agents. In this paper, we mainly focus on adversarial training and explore how adversarial training can enhance the robustness of DRL.

3. Preliminary

3.1. Deep Q-learning network

Deep Q-learning is a classical reinforcement learning method, and it can be formulated as a Markov Decision Process (MDP), which is defined with a tuple $\langle S, A, \pi(\theta), R, \gamma, \mu_0 \rangle$ where S is the state space, A is the action space, $\pi(\theta)$ is the policy network, R is the reward function, and γ denotes discount factor indicating how much the immediate reward will affect the future reward and μ_0 is the initial state. The agent tries to learn a policy network to interact with the environment efficiently. And the DQN aims to find an optimal policy network by solving the Bellman optimal equation:

$$Q(s, a) = E[r + \gamma \max_{a'} Q^*(s', a')] \quad (2)$$

Where r denotes the reward of the agent, s' denotes the current state, and a' is the action space. The general way is to train the network by minimizing the temporal different loss (e.g., squared Bellman error). And the best action is selected with the maximum Q value.

3.2. Strategically timed selection

As it is the sequential task for the agent to interact with the environment, to reduce the attacking frequency, and increase the attacking effectiveness, lots of studies focus on how to select the critical point (frame) to deploy the attack. A more general way is to select by inner driving where the attack is only deployed when the agent strongly prefers a specific action. And it is usually decided by the predefined attack score. In the policy gradient-based method, the attack score is usually defined as in Eq. (3).

$$c(s_t) = \max_{a_t} \pi(s_t, a_t) - \min_{a_t} \pi(s_t, a_t) \quad (3)$$

where π is the policy network which maps state-action pairs to a probability in action space, same intuition is always deployed in value-based methods (e.g., DQN), where the attack score is defined in Eq. (4).

$$c(s_t) = \max_{a_t} \Phi(Q(s_t, a_t)) - \min_{a_t} \Phi(Q(s_t, a_t)) \quad (4)$$

where Q is the value network, which Φ is the score function (e.g., softmax function). The attack is conducted when the attack score is larger than a threshold, which helps increase the attack effectiveness. While in some papers, to make the attack stealthy, they deployed the attack only when the variance of Q value (in Eq. (5)) is low.

$$v(s_t) = \frac{1}{|A| - 1} \sum_{i=1}^{|A|} (Q(s_t, a_i) - \frac{1}{|A|} \sum_{j=1}^{|A|} Q(s_t, a_j)) \quad (5)$$

Therefore, the attack rules are quite opposite when considering the different requirements. However, this type of attack depends on the threshold that is decided manually, and also they are lacking the comprehensive evaluation metric about the attacking effectiveness and stealthy.

3.3. Attacking methods

In this paper, the attack is deployed in a white-box environment where we know the parameters of the model, i.e., we need to know the actions of the target agent, as well as information about the state of each frame after the agent performs the action. Generally, for white-box attacks, the adversarial examples can be learned from as $x_{adv} = \operatorname{argmax}_{\|x' - x\|_\infty \leq \epsilon} l(h(x'), y)$, where x' denotes the crafted adversarial sample, and $l(\cdot)$ represents the loss function between the output and the correct label. There are several efficient attacking approaches (ATK for short) to craft the adversarial examples method, e.g., FGSM [22], Projected Gradient Descent (PGD) [30], and C&W [31], etc. Our paper also applied those three models to generate the adversarial samples under our framework. FGSM generates the adversarial sample with one step perturbation along the gradient direction, and the step size is constrained to ϵ [22]. The process is described as:

$$x_{adv} = x + \epsilon \cdot \operatorname{sign}(\nabla_x l(h(x), y)) \quad (6)$$

PGD subdivides the process into several small steps. And after each step of perturbation, all the adversarial samples are projected to the ϵ -ball again [30]. The process is described as:

$$x_{adv}^t = \Pi_\epsilon(x^{t-1} + \beta \cdot \operatorname{sign}(\nabla_x l(h(x^{t-1}), y))) \quad (7)$$

where β is the step size, and $\Pi(\cdot)$ represents the projection function. Generally, PGD is the strongest first-order attack. One of another popular attack method is C&W. The key idea of this attack method is to minimize the distance between normal samples and adversarial samples and to maximize the probability of the misclassification at the same time [31]. And the attack process is described as follows:

$$d_n = \frac{1}{2}(\tanh(w_n) + 1) - x_n \quad (8)$$

$$\min_{w_n} \|d_n\| + c \cdot f\left(\frac{1}{2}(\tanh(w_n) + 1)\right)$$

where $f(x') = \max(\max\{Z(x')_i : i \neq t\} - Z(x')_t, -k)$ with $Z(x')_t$ represents the probability of the adversarial sample x' , k represents the confidence, and d_n is the distance between the normal sample and the adversarial sample. In this paper, to fully understand our proposed attack framework, we apply all three methods to the attack.

4. Method description

In this paper, we focus on the white box attack after DRL is deployed, and try to deploy the attack with an intelligent agent which is named attacker by learning a stable policy network θ^X . With policy network θ^X , the attacker will know when to attack by learning a certain attack pattern over observations. The structure of the proposed model is shown in Fig. 1. It is the man-in-the-middle format for the attacker when it deploys the attack on state s_t . Then the adversarial image will be sent to the agent directly, the reward r^X of the attacker is originated from the agent. Based on this setting, the attacker's action space can be $\{None, Attack\}$, noted that *None* represents attacker not taking any action, and *Attack* represents attacker adding perturbation on observation. As we have mentioned above, most methods deploy the attack with the critical point to decide when conducting the attack [13,15], and this method is only based on rules and does not consider the long term reward, and is hard to obtain a good attacking performance. As we can see from the action space, the agent will be more stealthy if it takes more *None* operations. On the contrary, it is not. Generally, the goal of the model is to decrease DRL's performance with the fewest *Attack* steps. However, it is difficult

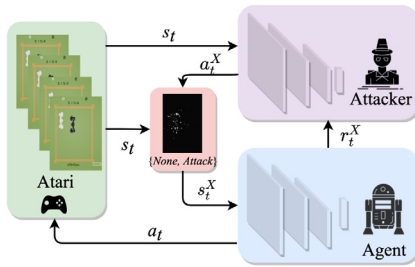


Fig. 1. The structure of the proposed model.

to make the balance between attack effectiveness and stealthy. Based on previous work [32], the goal of DRL is to maximize the expected value of the cumulative sum of a received scalar reward signal. Therefore, an effective way to guide the attacker is to design a combination of rewards, including effectiveness and stealth. To cope with this problem, we have designed two types of rewards to help the attacker learn a better policy network θ^X , and we will introduce them one by one.

4.1. Short-term attacking reward

Intuitively, if an agent is attacked and the attack is successful, its action will immediately deviate from its original action, otherwise, we will consider the attack a failure. To encourage a successful short-term attack, if the new action taken from the disturbing observation is different from the original action, we will give a positive reward to the agent. An important feature of this reward is that it changes over time. Moreover, the ultimate goal is to ensure that all attacks have a successful outcome. Therefore, we can use the success rate of the attack to represent the short-term return. In the long run, the success rate r_s of the attack can be defined as:

$$r_s = \left\lfloor \frac{N_a}{N_s} \right\rfloor \quad (9)$$

where $\lfloor \cdot \rfloor$ is the floor function to obtain one decimal place, N_a denotes the number of the successful attack before time t , and N_s is the total number of the attack before time t .

As mentioned above, it is best to implement stealth bonus in a short-term reward. This can be indicated by the attacking ratio. If the ratio of *None* is low, then the attacker has a low attack rate, and this can be considered an attack with good stealth. And the attack rate r_a is defined as:

$$r_a = \left\lfloor \frac{N_s}{N_t} \right\rfloor \quad (10)$$

where N_t denotes the total play of the agent before time t . From the work [33], we can conclude that an anomaly detector will be difficult to detect the attack when the attack number is small enough in deep reinforcement learning. Therefore, to ensure its stealth, we need to ensure the attacking number is low enough.

To balance the attack effectiveness and stealth, the reward function for attacker at time t is defined as:

$$r_{str} = r_s - r_a \quad (11)$$

This reward is also named a short-term reward because it shows the immediate effect of the agent performing the action. Different from previous work, we define the attack as a learning process where attackers need to learn the place (i.e., step) and the way (i.e., Attack on the action and attack on the state) to attack. Furthermore, suppose the attack success rate is 100% (i.e., $N_a = N_s$), and the agent play a oracle game which has $N_s = R_a$, then short-term reward will be $\frac{R_a - N_s}{R_a}$. From this function, we can conclude that the short-term reward is to ensure a minimal number of attacks by the attacker.

4.2. Long-term attacking reward

As we have described, there is no connection between the deviated action and its original action in the short-term attack, and the attacking action only depends on the current state, that is to say, there is no connection between the current attack and the future reward. However, to make up this problem, we need to find a good way to add the long-term effects to the attacker. One direct way is to apply the final rewards decrease to represent, and ensure the attack can accomplish some mission destruction tasks. Therefore, the long-term rewards is defined as followed:

$$r_{ltr} = \left[\frac{R_u - r}{R_u - R_l} \right] \quad (12)$$

If there are long-term effects over the agent, its final rewards will have a large decrease which can be represented by the percentage decline of the rewards. r denotes the accumulate reward of the current step $r = \sum_i^t r_i$, where r_i is the normal reward in the original setting. R_u is the upper bound of the accumulate reward, and R_l denotes the lower bound the accumulate reward. In our framework, the final reward is defined as a summation format with different ratio, i.e.,

$$r_t^X = \alpha r_{str} + (1 - \alpha) r_{ltr} \quad (13)$$

Where α is the hyperparameter that defines the ratio of r_{str} and r_{ltr} .

In general, the short-term reward is designed to stimulate agents to achieve attacks with fewer attacks which in turn ensures the stealthiness of the attacks, while the long-term reward is designed to ensure the effectiveness of the attacks by minimizing the accumulated rewards.

And the way to train the attacker X is designed as shown in Algorithm 1 by freezing the well-trained agent T . In the algorithm, M is the number of epochs, D is the replay memory of the attacker, and its capacity is set to N .

In the algorithm, the whole process is represented by lines 2–26, where in line 2 we need to initialize the agent state s_t at the beginning of each episode. Line 4 is the action selection of the trained agent T . Lines 5–10 indicate action selection by the attacker X using the *epsilon-greedy* strategy. Lines 11–19 represent the attacking steps in the training, where lines 11–12 represent that the attacker will do nothing when the action is *None*, lines 13–19 represent the procedures when action is not *None*, where line 14 is the attack with the *ATK* method, and when the attack is completed, a new state will be obtained and on which the agent will have a new action, which is in line 15. And lines 16–18 are to determine whether the attack is successful. Line 20 is to obtain the reward after the attack with Eq. (13), and line 21 is to store the transition sample into memory D for rollout training. Lines 22–24 indicate the model optimization procedures. Line 25 indicates updating the target action-value function \hat{Q} with Q every C steps.

5. Experiments

In the experiment, we will first verify the effectiveness of the proposed framework in the attack, then show the transferability of the attacking framework in different scenarios, i.e., transfer across the environment and transfer across the attacking methods, and finally apply adversarial training to enhance the robustness of DRL.


```

Input: Environment Atari, Attacker  $X$ , Agent  $T$ ;
Initialize replay memory  $D$  to capacity  $N$ ;
Initialize  $X$ 's parameters  $Q(s^X, a^X)$  with random weight  $\theta^X$  and
load the well-trained  $T$ 's weight  $\theta^T$ ;
Initialization target action-value function  $\hat{Q}$  of attacker with
 $\theta^X = \theta^T$ ;
1 for episode=1, $M$  do
2   Initialize status  $s_t$ ;
3   while not Done do
4     Select Agent  $T$ 's action  $a_t$ ;
5     Sample a value  $v$  from a uniform distribution  $[0, 1]$ ;
6     if  $v \leq \epsilon$  then
7       | Select a random action  $a_t^X$  from attack action space;
8     else
9       | select  $a_t^X = \operatorname{argmax}_{a^X} Q(s_t^X, a_t^X; \theta^X)$ ;
10    end
11    if  $a_t^X = \text{None}$  then
12      | Set  $s_t^X = s_t$ ;
13    else
14      | Obtained  $s_t^X$  with attack methods  $ATK$ ;
15      | Select new action  $a_{t_{new}} = \operatorname{argmax}_{a^X} Q(s_t^X, a_t; \theta^T)$ ;
16      | if  $a_{t_{new}} \neq a_t$  then
17        | Attack Success;
18      | end
19    end
20    Obtain  $r_t^X$  with Eq. (13);
21    Store transition  $(s_t^X, a_t^X, r_t^X, s_{t+1}^X)$  in  $D$ ;
22    Sample random minibatch of transition  $(s_j, a_j, r_j, s_{j+1})$ 
    from  $D$ ;
23    Set
24    
$$y_t = \begin{cases} r_j & \text{terminates at } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(s_{j+1}, a'; \theta^X) & \text{otherwise} \end{cases};$$

25    Perform a gradient descent step on  $(y_j - Q(s_j, a_j; \theta^X))^2$ 
    w.r.t. the parameter  $\theta^X$ ;
26    Set  $s_{t+1} = s_t$ ;
27    Every  $C$  steps reset  $\hat{Q} = Q$ ;
28 end

```

Algorithm 1: The attacking against the reinforcement learning with man-in-the-middle.

5.1. Settings

All of our experiments are run on the gym platform. The games selected include {Pong, Fishing Derby (Fishing for short), Boxing, Seaquest, MsPacMan (PacMan for short)}. The DRL is implemented in python based on Pytorch package,¹ and the input of Agent and the attacker from Atari's games are selected image frame which is transferred to 84×84 -pixel image, and the policy network is a classical convolutional neural network which is to map the input to the action space. There are three convolutional layers with the size of (32, 8, 8, 4), (64, 4, 4, 2) and (64, 3, 3, 1), where the first value in the bracket is the number of the filters, the second and the third values denote the filter size, and the last value represents the stride size. The last two layers are the fully connected functions that map the hidden representation to the action space, and the shape of the weights in these two layers are (3136, 512) and (512, action space). As we have described earlier, there are two actions for the attacker which are {None, Attack}. The activation function is ReLU. And there is the same structure of the policy neural network in Attacker. The structure of the policy network is illustrated in Fig. 2.

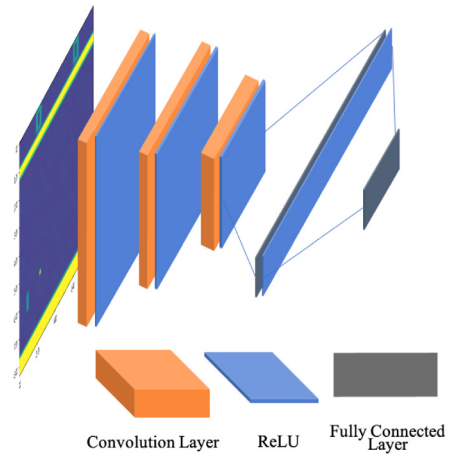


Fig. 2. The structure of the policy network.

The detailed hyper-parameters are listed in Table 1. Specifically, the learning rate is $1e-4$, and the batch size is 32. The replay memory buffer size is set to $1e6$, and the discount factor γ is 0.99. During the training, the current network updating frequency is 4, and the target network updating frequency is 100 by copying the parameters of $Q(s^X, a^X)$ to the target network (i.e., $C = 100$). In the ϵ -greedy policy, the start value is 1.0, and the end value is 0.01 with the fraction of number steps 0.1. The training starts after $1e4$ steps, and the maximum training step is $1e7$. To balance r_{str} and r_{ltr} , α is set to 0.5 in Eq. (13).

The evaluation metrics include the designed long-term reward r_{ltr} and short-term reward r_{str} . To make fair comparisons, all the experiments are sharing the same settings, and all the results are averaged from three times test. We will report the performance of different attacking methods, and the detailed models are listed as followed:

- **Uniform** Uniform denotes attacks are deployed in a uniform distribution. And this was proposed by Huang et al. [11].
- **Strategically Timed** Strategical Timed means the attack only happens at the point where the attack score is higher than the threshold, attack score can be calculated as in Eq. (4). And this was proposed by Lin et al. [13].

5.2. Attack evaluation

To evaluate the attacking effectiveness, we first train an oracle agent T on the games we selected. Then train the attacker X by freezing T . In order to have a good understanding of the attacking stealthy and effectiveness, all the results are reported in the designed balanced rewards r^X that are calculated in Eq. (13). And the results are shown in Table 2.

From the table, we can see that all the cases in our proposed framework have a better performance as it has higher attacking rewards. To our surprisingly, C&W is not the best-attacking model, on the contrary, and its performance is the worst one compared with the other two attacking approaches. One reason we think is that the attacking agent cannot learn a stable policy when facing the small state variance, as one of the objectives of C&W is to minimize the distance between the adversarial sample and the original sample as such it is a model-specific attack algorithm, therefore its use here makes its attack performance limited. To validate this hypothesis, we compared images that were attacked by different models at the same time, and these results are shown in Fig. 4. In the meantime, we also reported the time usage with different attacking approaches, and the results

¹ <https://pytorch.org/>

Table 1
The hyper-parameters applied in this paper.

Parameter	Value	Parameter	Value
Learning rate	1e-4	Batch size	32
Replay buffer D	1e6	Update current network frequency	4
Discount factor γ	0.99	Update target network frequency C	100
ϵ -greedy start threshold	1.0	Train stage	>1e4
ϵ -greedy end threshold	0.01	Fraction of number steps	0.1
α	0.5	Maximum number steps	1e7

Table 2
The attack evaluation with different attacking models, the higher the better, the value is calculated with Eq. (13).

Methods	Pong	Seaquest	Boxing	PacMan	Fishing
Uniform (FGSM)	0.77	0.67	0.66	0.64	0.70
Uniform (PGD)	0.81	0.64	0.70	0.67	0.70
Uniform (C&W)	0.04	0.54	0.37	0.56	0.54
Strategically Timed (FGSM)	0.86	0.74	0.79	0.59	0.62
Strategically Timed (PGD)	0.87	0.74	0.89	0.65	0.66
Strategically Timed (C&W)	0.07	0.62	0.46	0.40	0.47
Our (FGSM)	0.85	0.80	0.86	0.78	0.72
Our (PGD)	0.98	0.81	0.92	0.82	0.75
Our (C&W)	0.15	0.69	0.51	0.56	0.54

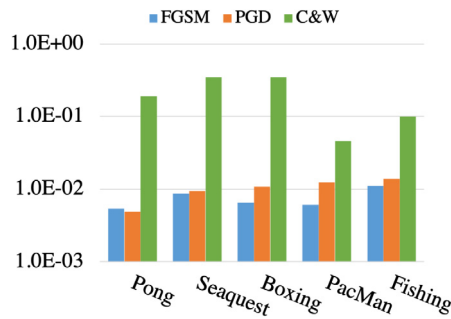


Fig. 3. The time with different attack approaches.

are shown in Fig. 3. From Fig. 4, we can see that the adversarial sample from C&W is the cleanest, and from Fig. 3, C&W has the longest time for attack conducting. And these make it difficult for the attacker to learn the adversarial strategy because of the long-time adversarial sample generation and the little change of the adversarial sample generated. However, under our proposed framework, PGD achieves the best performance. And this is also validated in the cases of “Strategically Timed” and “Uniform”.

5.2.1. Short-term reward & long-term reward

To make the results more clear, we use both short-term reward and long-term reward to demonstrate the attack reward r^x . The short-term attack denotes the single attack affects the agent’s behavior directly, i.e., high attacking success rate with the low attacking rate. And the long-term attack is the embodiment of the final effect of the series of attacks. And the results are shown in Fig. 5, and the points on the top right will have better performance in both immediate and long-term attacks.

In the figure, different attack strategies are shown in different shapes. ● denotes the “Uniform” attack, ■ is the “Strategically Timed” attack, and ◆ is our proposed attacking framework. And points of the same color represent the same type of attack: blue is the FGSM, brown is the PGD and green is the C&W. The straight lines are the contours that have the same r^x . From the figure, we can see that our proposed framework has the best performance as all nodes of the same type are below these lines. And especially

Table 3
The comparisons between different evaluation metrics.

Case	Rewards	r_s	r_a	r_{str}	r_{ltr}	r
1	-9.38	100%	56%	0.44	0.72	0.58
2	-8.69	100%	15%	0.85	0.71	0.78
3	-11.44	98%	4%	0.94	0.77	0.86

in the game of Pong, the attack with PGD under our proposed framework has the best performance by achieving the highest score in long-term reward and short-term reward at the same time. And this means our proposed framework has obtained the ability of effective stealthy attack. Specifically, there is a negative correlation between short-term reward and long-term reward, and the relationship between the two is difficult to balance, and this phenomenon is validated both in the cases of “Strategically Timed” and “Uniform”. Another interesting discovery in these two attacking scenarios is that all the games prefer having higher long-term rewards no matter what short-term rewards are, especially when the long-term reward is higher than the threshold of 0.5.

5.2.2. Advantages of the metrics

In this paper, we applied the newly designed metrics to evaluate the attack effectiveness and stealthy, to validate the wealth and goodness of the proposed new metrics, a comparison is made over the Pong game whose upper bound of the reward is 21, and the results are illustrated in Table 3. The general way to show the effectiveness of an attack is to use the reduction of the reward, which can be seen in the first column of Table 3. This approach is straightforward and clear in terms of attack effectiveness measurement and is used in many works [17,33]. However, we first need to know the upper bound of the reward and then can know how much the attack is reduced, meanwhile, this measurement does not reveal how many attacks have been launched against the agent, for example, in cases 1 and 2, the attack success ratio r_s against the agent is 100% in both cases, but the attack ratio r_a differs significantly. However, among our evaluation metrics, there are attack success ratio r_s and attack ratio r_a , which directly yield the short-term reward r_{str} , and this is a good metric to evaluate stealthiness. Therefore, in terms of our last column of Table 3, our proposed metric r balances effectiveness and stealthy well.

5.3. Transferabilities

How does the well-trained attacker behave when transferred to another environment, and how does the well-trained policy performs when we cooperate with another ATK (i.e., FGSM, PGD, and C&W). To answer these two questions well, we first transfer the well-trained attacker to different environments (i.e., games), and the details are presented in Section 5.3.1. We then cooperate the well-trained policy with different attacking methods, and the details are described in Section 5.3.2.

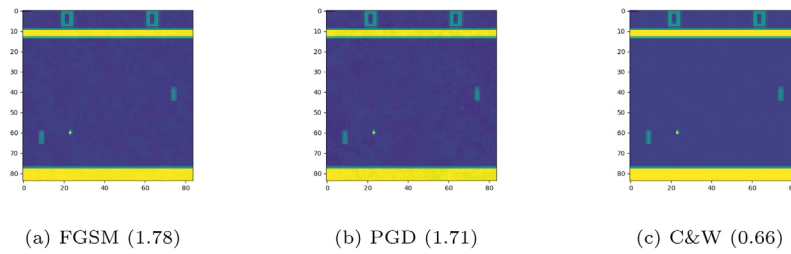


Fig. 4. The craft adversarial frame with different attack approaches, the value in the bracket indicates the absolute distance from the original frame.

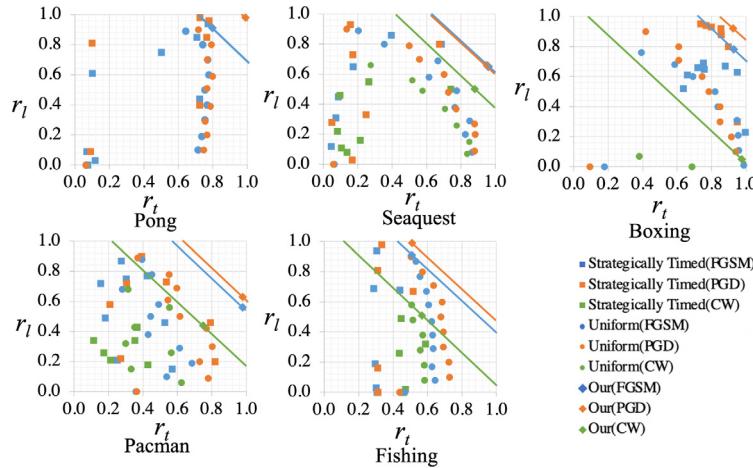


Fig. 5. The detailed results.

Table 4

Rewards r^x based on different transfer cases, “From” represents the original environment in which the attack model was trained, and “To” represents the game to which the attack model was transferred. The background color indicates the degree of the decrease compared to the normal condition. The darker the color, the greater the degree.

		From				
		Pong	Sequest	Boxing	PacMan	Fishing
To	Pong	0.98	0.04	0.37	0.54	0.04
	Sequest	0.48	0.81	0.51	0.52	0.48
	Boxing	0.47	0.10	0.92	0.60	0.10
	PacMan	0.51	0.48	0.50	0.82	0.48
	Fishing	0.64	0.15	0.29	0.35	0.75

Table 5

Rewards r^x based on different transfer cases, “Attacking Methods” are selected from FGSM, PGD, C&W, and “Policy” is the well-trained policy with different attacking methods. background color indicates the degree of the increase and decrease compared to normal condition, with green indicating the increase rate and red indicating the decrease rate. The darker the color, the greater the degree.

		ATK		
		FGSM	PGD	C&W
Policy	FGSM	0.80	0.76	0.64
	PGD	0.81	0.81	0.72
	C&W	0.75	0.71	0.69

5.3.1. Transferability across environments

In this scenario, we select PGD which performs the best as the attacking method. Then we transfer the well-trained policy to other environments, and the rewards after attacking are reported in Table 4.

From the table, we can see that almost all the attackers have a performance decrease compared with their original environment, and the performance of some cases drops as much as 90%. The worst case is the policy learned from Seaquest and Fishing, and they can hardly adapt in the new environment as more than half of the cases are in low reward r^x . However, there is still some transferring capability in our model. For example, with strategies learned from Pong, Boxing, and PacMan, the performance across different environments remains around 0.5. Based on those observations, we can see that when the attacker enters a new environment, it still needs to be fine-tuned or retrained to adapt to the new state.

5.3.2. Transferability across attacking methods

Generally, there are three well-trained policies under three attacking methods. To validate the transferability of the well-trained policy across different attacking methods and the attacking methods across different well-trained policies, we select the Seaquest as the test-bed, and we change different attacking methods described in line 14 in Algorithm 1. And the experiment results are reported in Table 5.

In the Table 5, with the FGSM learning the policy, the performance degradation degree increases successively when the attacking method is changed to PGD and C&W, and when using the PGD learning the policy, changing the attacking method to C&W will also degrade the performance, while changing the attacking method to FGSM does not. However, when the attacking method is changed to PGD and FGSM, the performance is improved successively by using the C&W learning strategy. From

these observations, it can be seen that the transferability of the attacking policy across different attacking methods is not as sensitive as the environment. The same applies to the transferability of the attacking methods across the different attacking policies. Compared with the simple attacking method (e.g., FGSM), the policy learned from a more complex attacking method (e.g., C&W) has better transferability. And the attack method PGD also has competitive performance compared with C&W, which is one of the reasons why we can get the best performance by using PGD in our framework.

5.4. Adversarial training

To validate the robustness of the model, we obtain the robust model (i.e., T_R) by fine-tuning the oracle DQN T with the well-trained attacker X that attacks with PGD, then validate the robustness with different attacking methods. The procedures are shown in Algorithm 2. And the experiment results are listed in Table 6.

Input: Environment Atari

- 1 Train an oracle agent T .
- 2 Freeze agent T , training attacker X with PGD.
- 3 Freeze the well-trained attacker X , fine-tune oracle agent T with adversarial samples to obtain the robustness model T_R .

Algorithm 2: The adversarial training procedure.

From the table, we can see that our adversarial training has boosted the model performance, and the improvement is ranging from 2.32% to 18.90%. That means adversarial training with the adversarial samples compensates for the vulnerability of the model to some degree. When we attack the robust agent T_R again, we can see that there is still a performance decrease in most of the cases, but the degradation is less than the attack that deployed directly on the Oracle agent T . In particular, the performance degradation of the PGD attack is the lowest, followed by the FGSM attack, and the performance degradation of the C&W attack is the largest. From these observations, we can see that the robustness of the model has indeed been improved to a certain extent after adversarial training, and the reason why PGD attack has the lowest performance degradation is that we used it to train the adversarial policy.

6. Conclusions & future works

In this paper, we proposed an intelligent attacker with reinforcement learning by considering the short-term attack, and long-term attack at the same time, then provided a useful evaluation metric to measure the attacking effectiveness and stealth. In our paper, the attack was deployed in a white-box environment where we know the actions of the target agent, as well as information about the state of each frame after the agent, performs the action, and this prior information can be easily obtained since the state and action are explicit information of the target agent. From the experimental results, we can see that the proposed attack method can achieve the goal in short steps attack. In the following step, we also conducted adversarial training to improve the robustness of reinforcement learning. In future work, we will explore more intelligent attack agents and try to explain why such attacks are useful in black-box scenarios, where one can employ parallel agents as target agents to accomplish policy network learning for the attack. Furthermore, we will explore how the robustness of the model under adversarial training and uncover the impact on the overestimate which is the inherent shortcoming of the DQN when deploying attacks.

Table 6

Adversarial training with the well-trained adversarial attack. The background color indicates the degree of the increase and decrease compared to normal condition, with green indicating the increase rate and red indicating the decrease rate. The darker the color, the greater the degree.

	Pong	Fishing	Boxing	Seaquest	PacMan
T	17.91	30.63	69.56	1177.50	804.69
T^R	19.87	31.34	74.66	1400.00	1020.00
Attack (FGSM)	18.59	27.22	40.65	980.00	587.19
Attack (PGD)	19.94	30.06	58.56	1080.00	660.00
Attack (C&W)	19.24	10.31	36.34	950.00	464.06

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research is supported by the National Natural Science Foundation of China (No. 62103330), and the Fundamental Research Funds for the Central Universities of China (3102021ZD-HQD09).

References

- [1] Yang Li, Wei Zhao, Erik Cambria, Suhang Wang, Steffen Eger, Graph routing between capsules, *Neural Netw.* 143 (2021) 345–354.
- [2] Haoran Li, Qichao Zhang, Dongbin Zhao, Deep reinforcement learning-based automatic exploration for navigation in unknown environment, *IEEE Trans. Neural Netw. Learn. Syst.* 31 (6) (2019) 2064–2076.
- [3] Yang Li, Quan Pan, Suhang Wang, Tao Yang, Erik Cambria, A generative model for category text generation, *Inform. Sci.* 450 (2018) 301–315.
- [4] Haiyun Peng, Yukun Ma, Soujanya Poria, Yang Li, Erik Cambria, Phonetic-enriched text representation for Chinese sentiment analysis with reinforcement learning, *Inf. Fusion* 70 (2021) 88–99.
- [5] Ke Zhang, Yuanqing Li, Jingyu Wang, Erik Cambria, Xuelong Li, Real-time video emotion recognition based on reinforcement learning and domain knowledge, *IEEE Trans. Circuits Syst. Video Technol.* (2021).
- [6] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, Rob Fergus, Intriguing properties of neural networks, in: 2nd International Conference on Learning Representations, ICLR 2014, 2014.
- [7] JinYin Chen, Yan Zhang, XueKe Wang, HongBin Cai, Jue Wang, ShouLing Ji, A survey of attack, defense and related security analysis for deep reinforcement learning, *Acta Automat. Sinica* 45 (2020) 1–19.
- [8] Edgar Tretschk, Seong Joon Oh, Mario Fritz, Sequential attacks on agents for long-term adversarial goals, 2018, arXiv preprint [arXiv:1805.12487](https://arxiv.org/abs/1805.12487).
- [9] Tong Chen, Jiqiang Liu, Yingxiao Xiang, Wenjia Niu, Endong Tong, Zhen Han, Adversarial attack and defense in reinforcement learning-from AI security view, *Cybersecurity* 2 (1) (2019) 1–22.
- [10] Xian Wu, Wenbo Guo, Hua Wei, Xinyu Xing, Adversarial policy training against deep reinforcement learning, in: 30th {USENIX} Security Symposium, {USENIX} Security 21, 2021.
- [11] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, Pieter Abbeel, Adversarial attacks on neural network policies, 2017, arXiv preprint [arXiv:1702.02284](https://arxiv.org/abs/1702.02284).
- [12] Vahid Behzadan, Arslan Munir, Vulnerability of deep reinforcement learning to policy induction attacks, in: *International Conference on Machine Learning and Data Mining in Pattern Recognition*, Springer, 2017, pp. 262–275.
- [13] Yen-Chen Lin, Zhang-Wei Hong, Yuan-Hong Liao, Meng-Li Shih, Ming-Yu Liu, Min Sun, Tactics of adversarial attack on deep reinforcement learning agents, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017, pp. 3756–3762.
- [14] Nicholas Carlini, David Wagner, Magnet and" efficient defenses against adversarial attacks" are not robust to adversarial examples, 2017, arXiv preprint [arXiv:1711.08478](https://arxiv.org/abs/1711.08478).
- [15] Jernej Kos, Dawn Song, Delving into adversarial attacks on deep policies, 2017, arXiv preprint [arXiv:1705.06452](https://arxiv.org/abs/1705.06452).
- [16] Alessio Russo, Alexandre Proutiere, Optimal attacks on reinforcement learning policies, 2019, arXiv preprint [arXiv:1907.13548](https://arxiv.org/abs/1907.13548).

- [17] Yiren Zhao, Ilia Shumailov, Han Cui, Xitong Gao, Robert Mullins, Ross Anderson, Blackbox attacks on reinforcement learning agents using approximated temporal information, in: 2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops, DSN-W, IEEE, 2020, pp. 16–24.
- [18] Chaowei Xiao, Xinlei Pan, Warren He, Jian Peng, Mingjie Sun, Jinfeng Yi, Mingyan Liu, Bo Li, Dawn Song, Characterizing attacks on deep reinforcement learning, 2019, arXiv preprint [arXiv:1907.09470](https://arxiv.org/abs/1907.09470).
- [19] Yi Han, Benjamin IP Rubinstein, Tamas Abraham, Tansu Alpcan, Olivier De Vel, Sarah Erfani, David Hubczenko, Christopher Leckie, Paul Montague, Reinforcement learning for autonomous defence in software-defined networking, in: International Conference on Decision and Game Theory for Security, Springer, 2018, pp. 145–165.
- [20] Panagiota Kiourti, Kacper Wardega, Susmit Jha, Wenchao Li, Trojdl: Trojan attacks on deep reinforcement learning agents, 2019, arXiv preprint [arXiv:1903.06638](https://arxiv.org/abs/1903.06638).
- [21] Xian Yeow Lee, Sambit Ghadai, Kai Liang Tan, Chinmay Hegde, Soumik Sarkar, Spatiotemporally constrained action space attacks on deep reinforcement learning agents, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, (04) 2020, pp. 4577–4584.
- [22] Ian J. Goodfellow, Jonathon Shlens, Christian Szegedy, Explaining and harnessing adversarial examples, 2014, arXiv preprint [arXiv:1412.6572](https://arxiv.org/abs/1412.6572).
- [23] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommannan, Girish Chowdhary, Robust deep reinforcement learning with adversarial attacks, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, 2018, pp. 2040–2042.
- [24] Vahid Behzadan, Arslan Munir, Whatever does not kill deep reinforcement learning, makes it stronger, 2017, arXiv preprint [arXiv:1712.09344](https://arxiv.org/abs/1712.09344).
- [25] Vahid Behzadan, William Hsu, Analysis and improvement of adversarial training in dqn agents with adversarially-guided exploration (age), 2019, arXiv preprint [arXiv:1906.01119](https://arxiv.org/abs/1906.01119).
- [26] Zhaoyuan Gu, Zhenzhong Jia, Howie Choset, Adversary a3c for robust reinforcement learning, 2019, arXiv preprint [arXiv:1912.00330](https://arxiv.org/abs/1912.00330).
- [27] Lerrel Pinto, James Davidson, Rahul Sukthankar, Abhinav Gupta, Robust adversarial reinforcement learning, in: International Conference on Machine Learning, PMLR, 2017, pp. 2817–2826.
- [28] Vahid Behzadan, Arslan Munir, Mitigation of policy manipulation attacks on deep q-networks with parameter-space noise, in: International Conference on Computer Safety, Reliability, and Security, Springer, 2018, pp. 406–417.
- [29] Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, Dmitry Vetrov, Variance networks: When expectation does not meet your expectations, in: International Conference on Learning Representations, 2018.
- [30] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, Adrian Vladu, Towards deep learning models resistant to adversarial attacks, in: International Conference on Learning Representations, 2018.
- [31] Nicholas Carlini, David Wagner, Towards evaluating the robustness of neural networks, in: 2017 IEEE Symposium on Security and Privacy, IEEE, 2017, pp. 39–57.
- [32] Richard S. Sutton, Andrew G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.
- [33] Jianwen Sun, Tianwei Zhang, Xiaofei Xie, Lei Ma, Yan Zheng, Kangjie Chen, Yang Liu, Stealthy and efficient adversarial attacks against deep reinforcement learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, (04) 2020, pp. 5883–5891.