

# New Challenges in Reinforcement Learning: A Survey of Security and Privacy

Yunjiao Lei<sup>1</sup>, Dayong Ye<sup>1</sup>, Sheng Shen<sup>1</sup> and Tianqing Zhu<sup>1\*</sup>

<sup>1\*</sup>School of Computer Science, University of Technology Sydney,  
Broadway, Sydney, 2007, NSW, Australia.

\*Corresponding author(s). E-mail(s): [Tianqing.Zhu@uts.edu.au](mailto:Tianqing.Zhu@uts.edu.au);  
Contributing authors: [Yunjiao.Lei@student.uts.edu.au](mailto:Yunjiao.Lei@student.uts.edu.au);  
[Dayong.Ye@uts.edu.au](mailto:Dayong.Ye@uts.edu.au); [Sheng.Shen-1@student.uts.edu.au](mailto:Sheng.Shen-1@student.uts.edu.au);

## Abstract

Reinforcement learning (RL) is one of the most important branches of AI. Due to its capacity for self-adaption and decision-making in dynamic environments, reinforcement learning has been widely applied in multiple areas, such as healthcare, data markets, autonomous driving, and robotics. However, some of these applications and systems have been shown to be vulnerable to security or privacy attacks, resulting in unreliable or unstable services. A large number of studies have focused on these security and privacy problems in reinforcement learning. However, few surveys have provided a systematic review and comparison of existing problems and state-of-the-art solutions to keep up with the pace of emerging threats. Accordingly, we herein present such a comprehensive review to explain and summarize the challenges associated with security and privacy in reinforcement learning from a new perspective, namely that of the Markov Decision Process (MDP). In this survey, we first introduce the key concepts related to this area. Next, we cover the security and privacy issues linked to the state, action, environment, and reward function of the MDP process, respectively. We further highlight the special characteristics of security and privacy methodologies related to reinforcement learning. Finally, we discuss the possible future research directions within this area.


**Keywords:** Reinforcement Learning, Security, Privacy Preservation, Markov Decision Process, Multi-agent System

# 1 Introduction

Reinforcement learning (RL) is one of the most important branches of AI. Due to its strong capacity for self-adaptation, reinforcement learning has been widely applied in multiple areas, including health care [1], financial markets [2], and robotics [3]. Reinforcement learning is considered to be a form of adaptive (or approximate) dynamic programming [4], and has achieved outstanding performance in solving complex sequential decision-making problems. Reinforcement learning's strong performance has led to its implementation and deployment across a broad range of fields in recent years, such as the Internet of things (IoT) [5], recommend systems [6], healthcare [7], robotics [8], finance [9], self-driving cars [10], and smart grids [11], and so on. Unlike other machine learning techniques, Reinforcement learning has a strong ability to learn by trial and error in dynamic and complex environments. In particular, it can learn from the environment which has minimum information about the parameters to be learned [12].

In the reinforcement learning context, an agent can be viewed as a self-contained, concurrently executing thread of control [13]. It can interact with the environment and obtain a state of the environment as input. The state of the environment can be the situation surrounding the agent's location. Take the road conditions in an autonomous driving scenario as an example. In figure 1, the green vehicle is an agent, and all the objects around it can be regarded as the environment; thus, the environment comprises the road, the traffic signs, other cars, etc. Based on the state of the environment, the agent chooses an action as output. Next, the action changes the state of the environment, and the agent will receive a scalar signal that can be regarded as an indicator of the value for the state transition from the environment. This scalar signal is always represented as a reward. The agent's purpose is to learn an optimal policy over time by trial and error in order to gain a maximal accumulated reward as reinforcement. In addition, the combination of deep learning and reinforcement learning further enhances the ability of reinforcement learning [14].

However, reinforcement learning has been subject to privacy attacks due to its weaknesses that can be leveraged by attackers. Established samples used in reinforcement learning contain the learning agent's private information, which is vulnerable to a wide variety of attacks. For example, in disease treatment applications with reinforcement learning [1], real-time health data is required, and to achieve an accurate dosage of medicine, the information is always collected and transmitted in plaintext. This may cause disclosure of users' private information; consequently, the reinforcement learning system may collect data from public resources. Most collected datasets contain private or sensitive information that has a high probability of being disclosed [15]. Moreover, reinforcement learning may also require data sharing [16] and needs to transmit information during the sharing process. Thus, attacks on network links can also be successful in a reinforcement learning context. Furthermore, cloud computing, which is always used for reinforcement learning computation and



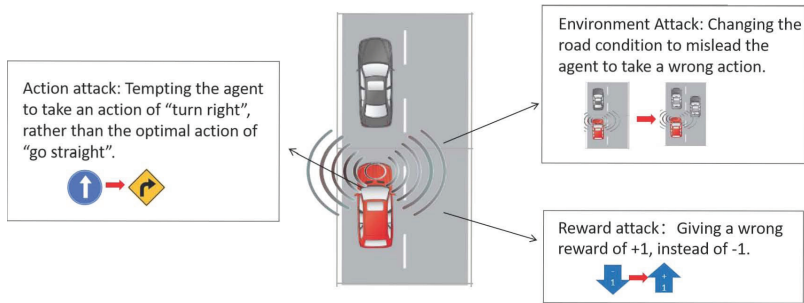
**Fig. 1** An autonomous driving scenario. The green car is an agent. the environment comprises the road, the traffic signs, other cars, etc.

storage has inherent vulnerabilities to certain attacks [17]. Rather than changing or affecting the model, the attackers may choose to focus on obtaining or inferring the privacy data; for example, Pan et al. [18] inferred information about the surrounding environment based on transition dynamics.

Moreover, reinforcement learning is also weak to security attacks. It is tender for attackers to leverage the breachable data source [19]. For example, data poisoning attacks [20] and adversarial perturbations [21] are very popular existing approaches in this field. From a defense perspective, several methods have been proposed over the past few years to address these security concerns. Some researchers have focused on protecting the model from attacks and ensuring that the model still performs well while under attack. The aim is to make sure the model takes safe actions that are exactly known, or to get optimal policy under worse situations, such as by using adversarial training [22].

Figure 2 presents an example of security attacks in reinforcement learning in an autonomous driving scenario. An autonomous car is driving on the road and observing its environment through sensors. To keep safe while driving autonomously, it will continually adjust its behavior based on the road conditions. In this case, an attacker may focus on influencing the autonomous driving conditions. For example, at a particular time, the optimal action for the car to take is to go straight; however, an action attack may directly influence the agent to turn right(attack may also impact the value of the reward). With regard to environmental influencing attacks, the attacker may conceive or falsely insert a car in the right front of the environment, and this disturbing may mislead the autonomous car into taking a wrong action. As for reward attacks, rivals may try to change the value of the reward(e.g., from +1 to -1) and thereby impact the policy of the autonomous car.

The main approaches to defending privacy and security in the reinforcement learning context include encryption technology [23] and information-hiding techniques, such as differential privacy [24]. In addition, some artificial algorithms also have been used to preserve individual privacy [25], such as federated



**Fig. 2** A simple example of a security attack in reinforcement learning in the context of automatic driving. An action attack, environmental attack and reward attack are shown respectively. An action attack works by influencing the choice of action directly, such as by tempting the agent to take the action "turn right" rather than the optimal action "go straight". Environmental attacks attempt to change the agent's perception of the environment so as to mislead it into taking an incorrect action. Finally, the reward attack works by changing the value of a reward given for a specific action in a state.

learning (FL) which can preserve privacy for the learning mechanism and structure. Yu et al. [26] adopt federated learning (FL) into a deep reinforcement learning model in a distributed manner, with the goal of protecting data privacy for edge devices.

As an increasing number of security and privacy issues in reinforcement learning emerge, it is meaningful to analyze and compare existing studies to help spark ideas about how security and privacy might be improved in future in this specific field. Chen et al. [27] reviewed the research related to reinforcement learning from the perspective of artificial intelligence security about adversarial attacks and defence. Another survey paper [12] conducted a literature review on securing IoT devices using reinforcement learning. Two further survey papers [28], [29] explored deep reinforcement learning in IoT.

However, the works mentioned above are all focused on the IoT or communication networks. Very few existing surveys have comprehensively presented the security and privacy issues in reinforcement learning. Accordingly, in this paper, we highlight the objects that the attacks aim at and provide a comprehensive review of the key methods used to attack and defend these objects.

The main contributions of our survey can be summarized as follows:

- The survey organizes the relevant existing studies from a novel angle that is based on the components of the Markov decision process (MDP). We classify current researches on attacks and defences based on their objects in MDP. This provides a new perspective that enables focusing on the target of the methods across the entire learning process.
- The survey provides a clear account of the impact caused by the targeted objects. These objects are components in MDP that are related to each other and may exist in the same time or/and space. Adopting this approach

enables us to follow the MDP to comprehend the relevant objects and the relationships between them

- The survey compares the main methods of attacking or defending the components of MDP, and thereby sheds some lights on the advantages and disadvantages of these methods.

The remainder of this paper is structured as follows. We first present preliminary concepts in reinforcement learning systems in Section 2. We then outline the security and privacy challenges in reinforcement learning in Section 3. Next, we present further details on security in reinforcement learning in Section 4, followed by an overview of privacy in reinforcement learning in Section 5. Finally, Sections 6 and 7 present our avenues for future work and conclusion respectively.

## 2 Preliminary

### 2.1 Notation

Table 1 lists the notations used in this article. RL is reinforcement learning, and DRL is deep reinforcement learning. MDP stands for the Markov Decision Process, which is widely used in reinforcement learning. MDP can be denoted by a tuple  $(S, A, T, r, \gamma)$ , which is made up of the agent action space  $A$ , the environment state space  $S$ , the reward function  $r$ , the transition dynamics  $T$ , and a discount factor  $\gamma \in [0, 1)$ . The transition dynamics is a probability mapping from state-action pairs to states  $T : (S \times A) \times S \rightarrow [0, 1]$ . The agent's purpose is to find an optimal policy that can map environment states to agent actions to maximize long-term reward.  $v^\pi(s)$  and  $Q^\pi(s, a)$  are the state and action-state values, which can regard as a means of evaluating the policy.

**Table 1** The main notations through the paper.

notations	meaning
$RL$	Reinforcement learning
$DRL$	Deep reinforcement learning
$MDP$	Markov decision process
$A$	The action space of the agent
$S$	The state space of the environment
$T$	The transition dynamics
$r$	The reward function
$\gamma$	A discount factor which is within the range (0,1)
$\pi$	Policy
$v^\pi(s)$	State value
$Q^\pi(s, a)$	Action-state value

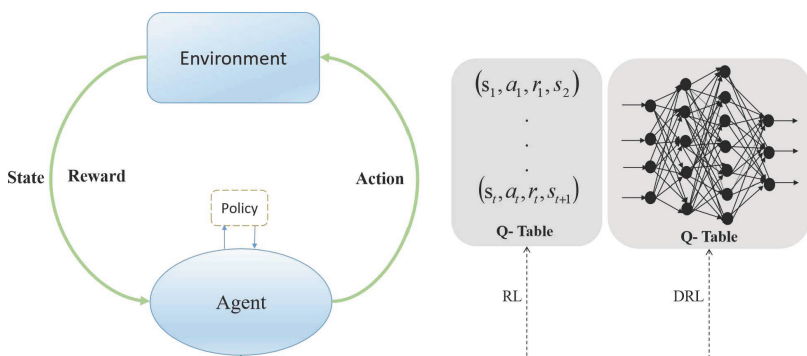
### 2.2 Reinforcement learning

The reinforcement learning model contains the environment states  $S$ , the agent actions  $A$ , and scalar reinforcement signals that can be regarded as rewards

$r$ . All the elements and the environment can be conceptualized as a whole system. At step  $t$ , when an agent interacts with the environment, it can receive a state of the environment  $s_t$  as input. Based on the state of the environment  $s_t$ , the agent chooses an action  $a_t$  using the policy  $\pi$  as output. Next, the action changes the state of the environment to  $s_{t+1}$ . At the same time, the agent will obtain a reward  $r_t$  from the environment. This reward is a scalar signal that can be regarded as an indicator of the value for the state transition.

In this process, the agent learns a piece of knowledge, which may be recorded as  $s_t, a_t, r_t, s_{t+1}$  in a Q table. Q table has calculated the maximum expected future rewards for action at each state, and can guide us to choose the best action at each state. In the next step, the updated  $s_{t+1}$  and  $r_{t+1}$  will be sent to the agent again. The agent's purpose is to learn an optimal policy  $\pi$  so as to gain the highest possible accumulated reward  $r$ . To arrive at the optimal policy  $\pi$ , the agent can train by applying a trial and error approach over the long-term episodes.

A Markov Decision Process (MDP) with delayed rewards is used to handle reinforcement learning problems, such that MDP is a key formalism in reinforcement learning.



**Fig. 3** The interaction between agent and environment with MDP. The agent interacts with the environment to gain knowledge, which may be recorded as a table or a neural network model (in DRL), and then takes an action that will react to the environment state.

If the environment model is given, two simple iterative algorithms can be chosen to arrive at an optimal model in the MDP context: namely, value iteration [30] and policy iteration [31]. When the information of the model is not known in advance, the agent needs to learn from the environment to obtain this data based on an appropriate algorithm, which is usually a kind of statistical algorithm. Adaptive Heuristic Critic and  $TD(\lambda)$ , which is a policy iteration mechanism, were used in the early stages of reinforcement learning to learn an optimal policy with samples from the real world [32]. Subsequently, the Q-learning algorithm increased in popularity [33, 34] and is now also a very important algorithm in reinforcement learning. The Q-learning algorithm is also an iterative approach used to select an action with a maximum Q value,

which is an evaluation value, in order to ensure that the chosen policy is optimal. Moreover, due to its ability to deal with high-dimensional data and to approximate the function, deep learning has been combined with reinforcement learning to create the field of “deep reinforcement learning” (DRL) [35]. This combination has led to significant achievements in several fields, such as learning from visual perceptual [14] and robotics [36].

An example of reinforcement learning is presented in Figure 4. The figure depicts a robot searching for an object in the Grid World environment. The red circle represents the target object, the grey boxes denote the obstacles, and the white boxes denote the road. The robot’s purpose is to find a route to the red circle. At each step, the robot has four choices of action: walking up, down, left and right. In the beginning, the agent receives information from the environment which may be obtained through sensors such as radar or camera. The agent then chooses an action and receives a corresponding reward. In the position shown in the figure, choosing the action of up, left or right, may result in a lower reward, as there are obstacles in these three directions. However, taking the action of moving down will result in a higher reward, as it will bring the agent closer to its goal.



**Fig. 4** A simple example of reinforcement learning, in which a robot tries to find an object in the Grid World environment. The blue robot can be seen as the agent in reinforcement learning. The red circle is the target object. The grey boxes denote the obstacles, while the white boxes denote the road. The robot’s purpose is to find a route to the red circle.

## 2.3 Markov Decision Process (MDP)

The Markov decision process (MDP) is a framework used to model decisions in an environment [37]. From the perspective of reinforcement learning, MDP is an approach which has a delayed reward. In MDP, the state transitions are not related to any previous environment states or agent actions. That is to say,

the next state is independent of the previous states and based on the current environment state.

MDP can be denoted as the tuple  $(S, A, T, r, \gamma)$ , which is made up of the agent action space  $A$ , the environment state space  $S$ , the reward function  $r$ , the transition dynamics  $T$ , and a discount factor  $\gamma \in [0, 1]$ . The transition dynamics can be defined as a probability mapping from state-action pairs to states  $T : (S \times A) \times S \rightarrow [0, 1]$ . The agent's purpose is to find an optimal policy  $\pi$  that can map environment states to agent actions in a way that maximizes its long-term reward. The discount factor  $\gamma$  is applied to the accumulated reward to discount future rewards. In many cases, the goal of a reinforcement learning algorithm with MDP is to maximize the expected discounted cumulative reward.

At time step  $t$ , we denote the environment state, agent action, and reward by  $s_t$ ,  $a_t$  and  $r_t$  respectively. Moreover, we use  $v^\pi(s)$  and  $Q^\pi(s, a)$  to evaluate the state and action-state value. The state value function can be expressed as follows:

$$v^\pi(s) = E_\pi \left[ \sum_{t \geq 0} \gamma^t r(s_t, a_t) | s_0 = s, \pi \right] \quad (1)$$

The action-state value function is as follows:

$$Q^\pi(s, a) = E_\pi \left[ \sum_{t \geq 0} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a, \pi \right] \quad (2)$$

In a wide variety of works, Q-learning was the most popular iteration method applied to discounted infinite-horizon MDPs.

## 2.4 Deep reinforcement learning

In some cases, reinforcement learning finds it difficult to deal with high-dimensional data, such as visual information. Deep learning enables reinforcement learning to address these problems. Deep learning is a type of machine learning that can use low-dimensional features to represent high-dimensional data through the application of a multi-layer Artificial Neural Network (ANN). Consequently, it can work with high-dimensional data in fields such as image and natural language processing. Moreover, deep reinforcement learning (DRL) combines reinforcement learning with deep neural networks, thereby enabling reinforcement learning to learn from high-dimensional situations. Hence, DRL can learn directly from raw, high-dimensional data, and can accordingly acquire the ability to understand the visual world. Moreover, DRL also has a powerful function approximation capacity, which also employs deep neural networks to train approximate functions in reinforcement learning; for example, to produce the approximate function of action-state value  $Q^\pi(s, a)$  and policy  $\pi$ .



The process of DRL is nearly the same as that of reinforcement learning. The agent's purpose is also to obtain an optimal policy that can map environment states to agent actions in a way that maximizes long-term reward. The main difference between the DRL and reinforcement learning processes lies in the Q table. As shown in Figure 3, in reinforcement learning, this table may be a form that records the map from state to action; by contrast, in deep reinforcement learning, a neural network is typically used to represent the Q table.

## 3 Security and privacy challenges in reinforcement learning

In this section, we will briefly discuss some representative attacks that cause security and privacy issues in reinforcement learning. In more detail, we explore different types of security attacks (specifically, adversarial and poisoning attacks) and privacy attacks (specifically, genetic algorithm (GA) and inverse reinforcement learning (IRL)). Moreover, some representative defence methods will also be discussed (specifically, differential privacy, cryptography, and adversarial learning). We further present the taxonomy based on the components of MDP in this section, along with the relationships and impacts among these components in reinforcement learning.

### 3.1 Attack methodology

#### 3.1.1 Security attacks

In this part, we discuss security attacks designed to influence or even destroy the reinforcement learning model in the reinforcement learning context. Specifically, we briefly introduce some recently proposed attack methods developed for this purpose.

One of the popular meanings of the term "security attack" is an adversarial attack with adversarial examples [38, 39]. The common form of adversarial examples involves adding imperceptible perturbations to data with a pre-defined goal; these perturbations can deceive the system into making mistakes that cause malfunctions, or prevent it from making optimal decisions. Because reinforcement learning gathers examples dynamically throughout the training process, attackers can directly add imperceptible perturbations to states, environment information, and rewards, all of which may influence the agent during reinforcement learning training. For example, consider the addition of tiny perturbations to state  $s$  in order to produce  $s + \delta$  [35, 40] ( $\delta$  is the added perturbation). Even this small change may affect the following reinforcement learning process. Attackers determine where and when to add perturbations, and what perturbations to add, in order to maximize the effectiveness of their attack.

Many algorithms that add adversarial perturbations have been proposed. Examples include the fast gradient sign method (FGSM), which can calculate

adversarial examples, the strategically-timed attack, which focuses on selecting the time step of adversarial attacks, and enchanting attack (EA), which can mislead the agent regarding the expected state through a series of crafted adversarial examples. Moreover, defenses to adversarial examples have also been studied. The most representative method is adversarial training [41], which trains agents under adversarial examples and thereby improves model robustness. Other defensive methods focus on modifying the objective function, such as by adding terms to the function or adopting a dynamic activation function.

Another common type of security attack is the poisoning attack, which focuses on manipulating the performance of a model by inserting maliciously crafted "poison data" into the training examples. A poisoning attack is often selected when an attacker has no ability to modify the training data itself; instead, the attacker adds examples to the training set, and those examples can also work at test time. Attacks based on a poisoned training set aim to influence the behaviour of the model so that it outputs incorrect results. As reinforcement learning requires a very large amount of data for training, and may also collect various types of data from sensors and public applications, it may be vulnerable to fake data injected by attackers into these kinds of data inputs. Examples include the poisoning attack on the environment [19], in which the attacker crafts malicious environmental examples (such as the transition dynamics) at training time to change the policy.

The most effective method of crafting poisoned data may be the traditional gradient-based method, and there are also some methods based on the gradient method. The representative methods to defend against poisoning attacks are detection methods and training-based defences. The detection method attempts to detect the poisoned training data or identify corrupted models after they have been trained. The training-based defences are designed to develop robust training routines or to remove the effects of poisoned data.

### 3.1.2 Privacy attacks

Two common types of privacy attacks are those that get/search private information directly and those that infer private information based on known information. Data transfer and storage are necessary components of reinforcement learning, as the agent requires a large amount of data for training and also needs to interact with its environment. As a consequence, privacy attacks on storage and transfer can be also used for reinforcement learning. Moreover, some other special inferring methods are also used.

Genetic algorithm (GA) belong to the category of evolutionary computing algorithms, which are a type of search algorithm inspired by the process of natural selection. These algorithms can be used to attack reinforcement learning systems in order to obtain privacy information. Transition dynamics search [18] is one such method. The basic Genetic Algorithm starts with a randomly initialized population of candidates. A selection operator selects parents by randomly picking two candidates and choosing the one with the

higher score. Crossover is then used to generate the child candidates of selected parents, and random mutation is applied to the child candidates to generate new candidates [42].

Inverse reinforcement learning (IRL) is a kind of inferring algorithm that can be used to infer the reward function of reinforcement learning based on the policy or observed behaviour [43]. In inverse reinforcement learning methods, the observed agent is regarded as an expert while the subject agent is viewed as the learner, and IRL assumes that the expert is behaving according to an underlying policy. The purpose of IRL is to learn an optimal reward function that can explain the observed behaviours. IRL is usually used to help the reinforcement learning system to obtain a reward function; however, because of its ability to infer, IRL may also be applied to attack.

## 3.2 Possible defense methodologies

In this section, we will present three representative defensive methods that are widely used in various fields.

### 3.2.1 Differential Privacy

Differential privacy is a prevalent privacy model that can guarantee minimal impact on the analytical output of a dataset if any individual record is stored in or removed from a dataset [44].

In differential privacy, two datasets  $D$  and  $D'$  are regarded as neighbouring datasets if they differ in terms of only one record. A query  $f$  is a function that maps the dataset  $D$  to an abstract range  $R$  ( $f : D \rightarrow R$ ). A group of queries is denoted as  $F = \{f_1, \dots, f_m\}$ , and  $F(D)$  denotes  $\{f_1(D), \dots, f_m(D)\}$ . The maximal difference in the results of query  $f$  is defined as the sensitivity of query  $\Delta f$ , which determines how much perturbation is required for a privacy-preserving answer at a given privacy level. The goal of differential privacy is to mask the differences in the answers to query  $f$  between the neighbouring datasets. To achieve this goal, differential privacy provides a mechanism  $M$ , which is a randomized algorithm that accesses the datasets.

There are three types of widely used differential privacy mechanisms: the Laplace mechanism, the exponential mechanism and the Gaussian mechanism. The Laplace mechanism adds Laplace noise to the true answer; here,  $\text{Lap}(b)$  is used to represent the noise sampled from the Laplace distribution with scaling  $b$ . Exponential mechanisms  $M$  select and output an element with probability proportionality. Compared to a Laplace mechanism, a Gaussian mechanism adds zero-mean isotropic Gaussian distribution sampled noise.

Differential privacy can be used in learning problems to improve various aspects of a model, such as randomization, privacy preservation capability, and algorithm stability [45]. For example, Ye et al. [46] applied differential privacy to a confidence score vector containing a probability distribution over the possible classes predicted by an ML model. This approach can defend against data inference attacks in a time-efficient manner by controlling the utility loss

of confidence score vectors; in so doing, it fully reflects the advantages of the privacy preservation capability and algorithm-stability ability of differential privacy.

There are two popular variants of differential privacy: joint differential privacy (JDP) and local differential privacy (LDP). The former has a centralized agent that is responsible for protecting users' sensitive data, while in the latter, information needs to be protected directly on the user side.

### 3.2.2 Cryptography

Cryptography is the classic method used for privacy protection fields by encoding messages so that they cannot be understood by untrusted parties. The main encoding techniques are symmetric algorithms and asymmetric algorithms. Symmetric algorithms utilize the same key for both encryption and decryption. Examples include Data Encryption Standard (DES), triple-DES (TDES) and Advanced Encryption Standard (AES). DES was the first encryption standard method. It employs a block cipher that can encrypt 64 bits of plain text at a time, along with a 56-bit key. The TDES algorithm adopts three rounds of DES encryption, so that it has a key length of 168 ( $56 * 3$ ) bits, along with two or three 56-bit keys. This method first uses three different keys to generate the cipher text  $C(t)$  from the plaintext message  $t$ . One of the most popular Cryptography methods in recent years is Homomorphic Cryptosystems [23], which allows operations on the cipher text; hence, it has great adaptability and is suitable for use in different systems for different aims [47]. A homomorphic encryption algorithm  $H$  is a set of four functions  $H = KeyGeneration, Encryption, Decryption, Evaluation$ . Here, key generation is a client that generates a pair of keys: a public key  $pk$  and a secret key  $sk$  for encryption of plain text. The purpose of Encryption is to use the  $sk$  client to encrypt the plain text PT and generate  $Esk(PT)$ . The cipher text (CT) will be sent to the server with the public key  $pk$ . The Evaluation function evaluates the cipher text (CT). Finally, the Decryption function uses  $sk$  to decrypt and obtains the original result.

### 3.2.3 Adversarial learning

Some smart methods have been adopted for privacy protection purposes. One of these methods, adversarial learning, is used specifically to combat adversarial attacks. The main idea behind adversarial learning involves training a model on a training set with adversarial examples to increase model robustness. To do this effectively, it may be necessary to generate a large number of adversarial examples or increase the amount of perturbed data. During training time, an agent may learn with a modified objective function that has the original loss function  $J$ . The objective function is as follows:

$$\tilde{J}(\theta, x, y) = \alpha J(\theta, x, y) + (1 - \alpha)J(\theta, x + \epsilon \text{sign}(\nabla_x J(\theta, x, y)), y) \quad (3)$$

Notably, these methods may require a huge number of adversarial examples, and may also be computationally intensive. Moreover, some of the defensive strategies do not work against certain kinds of adversarial attacks.

### 3.3 Taxonomy based on the components of MDP

In the below, we classify security and privacy in reinforcement learning based on the components of MDP. First, we categorize the existing research as either security-related or privacy-related. In this paper, we regard papers about attacks that influence or even destroy the reinforcement learning model and defences that enhance the robustness of the reinforcement learning model as examples of security problems in reinforcement learning. For example, studying how to attack a system to mislead the agent, or how to protect the system from adversarial attacks so that it remains stable and produces good output. In addition, research into obtaining, inferring, or conversely protecting private information will be regarded as related to privacy in reinforcement learning in this paper. Examples include inferring the information of the environment based on known transition dynamics and using cryptograms to preserve privacy.

We subsequently conduct a further classification of security in reinforcement learning and privacy in reinforcement learning based on the MDP perspective. In a reinforcement learning model, an agent interacts with the environment via perception and action, and a Markov Decision Process (MDP) is always used for reinforcement learning (as shown above). MDP consists of the tuple  $(S, A, T, r, \gamma)$ . We can thus organize the categorization following the Markov Decision Process, especially the elements of the MDP tuple. Specifically, we identify the attack and defense targets of state and action, environment and reward. The state  $s$  and action  $a$  refer to the state of the environment and the action of the agent in reinforcement learning, or an expression based on  $s$  or/and  $a$  (e.g., Q-value). The environment includes the transition dynamics in MDP and surrounding environment situations. The reward aspect pertains to studies that aim at the reward function of MDP in reinforcement learning. The taxonomy is shown in figure 2.

The agent's purpose is to find an optimal policy that can map environment states to agent actions. To facilitate more efficient policy evaluation, the concept of reward was introduced to reinforcement learning. A reward is a scalar signal that can be regarded as an indicator of the value for the state transition. Long-term rewards can be adopted to assess the policy; examples include the mean value of the reward, accumulated reward, or other functions based on rewards, such as the Q-value and V-value.

At each step, the agent first observes the state  $s_t$  from the environment. Next, the agent takes an action  $a_t$  based on the policy  $\pi$ ; subsequently, the agent will receive a reward  $r_t$  as feedback from the environment, and the environment changes to  $s_{t+1}$ . The transition dynamics can be defined as a probability mapping from state-action pairs to states  $T : (S \times A) \times S \rightarrow [0, 1]$ . It is usually difficult to determine an optimal policy based solely on

Table 2 Taxonomy.

Section	Subsection	Target	Impact
Security in reinforcement learning	Security of state and action in MDP	Action	State, Environment, Reward, Policy
	Security of environment in MDP	State	Action, State, Environment, Policy
	Security of reward function in MDP	Transition dynamics Surrounding situations	State, Reward, Policy State, Reward, Policy
Privacy in reinforcement learning	Privacy of state and action in MDP	Reward	Policy, Action, State
	Privacy of environment in MDP	Action	Policy
	Privacy of reward function in MDP	State	Action, State, Environment, Policy
	Privacy of state and action in MDP	Transition dynamics Surrounding situations	State, Action, Reward, Policy State


In this article, we first divide security and privacy in reinforcement learning into two main sections: security in reinforcement learning and privacy in reinforcement learning. We then further classify the related research from the perspective of MDP.

Every main section will have three subsections: state and action, environment, and reward. These three aspects are all potential attack/defense targets, and may have some impact on the process.

the calculations of the evaluating function. Hence, value iteration methods are adopted to solve this issue. The action-state value iteration formula for evaluating the policy can be expressed as follows:

$$Q_{i+1}(s, a) = Q_i(s, a) + \alpha \left[ r + \gamma \sum \pi(a_{t+1}|s_{t+1})Q_i(s_{t+1}, a_{t+1}) - Q_i(s_t, a_t) \right] \quad (4)$$

Figure 5 illustrates the process, along with the attacks of the main elements classified in this paper.



**Fig. 5** The objects of attacks/defenses and their impacts on the reinforcement learning process. We can observe that all elements are situated in the chain of the reinforcement learning process; they are not isolated, but connected to each other. Therefore, attacks aimed at  $s_t$ ,  $a_t$ ,  $r_t$ , and the environment (such as the transition dynamics  $T$ ) can all affect the elements in the chain.

We can determine that the state  $s_t$  of the environment may affect the action taken. In more detail, an agent using the policy output an action based on the state. The action  $a_t$  will in turn affect the environment, and may impact the reward  $r_t$  and the next state  $s_{t+1}$ . Moreover, the function used to evaluate the policy is based on  $s_t$ ,  $a_t$  and  $r_t$ . Consequently, if an adversary attacks  $s_t$ , the action  $a_t$  may be directly influenced, while the reward and policy may also be impacted. If opponents attack  $a_t$ , the next state and reward may be changed by this attack; furthermore, the policy may also be affected. In addition, attacks targeting rewards can also influence the policy, while attacks on the environment (like the transition dynamics, which is a probability mapping from state-action pairs to states) also have an influence on the choices of action and the next state. Thus, attacks aimed at  $s_t$ ,  $a_t$ ,  $r_t$ , and the environment (such as the transition dynamics  $T$ ) can all influence the agent to make a sub-optimal decision. In addition, the learning process is continuous; thus, attacks on every component in the chain may affect each other in some steps, and as a result, the agent may fail to achieve its goal.

## 4 Security in reinforcement learning

Security is one of the most significant aspects of reinforcement learning, as it involves exploring possible attacks and defenses to improve the robustness of the model. It tends to get a model which has a reliable and stable performance even when faced with sudden interference or changes.

In this section, we will present a review of security in reinforcement learning based on the MDP perspective. This section can be divided into three

subsections. The first addresses the security of state and action in MDP. The second examines the security of the environment, which includes the transition dynamics in MDP and the surrounding environment situations. Finally, we explore the security of the reward function in MDP.

## 4.1 Security of state and action in MDP

The state and action in MDP are the state of the environment and the action of the agent in reinforcement learning. Through the use of adversarial examples or the application of sudden perturbations to these elements, a reinforcement learning model can be misled from optimal performance. Many existing methods focus on attacking the state and action in reinforcement learning; at the same time, methods of defending against these attacks have also been studied.

In this section, we attempt to survey the security problems that target the state or/and action in MDP of reinforcement learning. Lee et al. [48] and Chen et al. [49] focused on the action in MDP and studied the attacking and protecting of actions respectively. Zhao et al. [40], Garrett et al. [50], Sun et al. [35] and Ye et al. [51] all focused on attacking states in MDP of reinforcement learning. Moreover, Dai et al. [52] learned to protect both the actions and states in MDP.

### 4.1.1 Action Security in MDP

Attackers targeting actions can misdirect the agent's subsequent progress, and may also influence the reward or policy. One of the popular basic approaches to such attacks involves the use of adversarial samples that add perturbations to the action space. It is also possible to disturb the action sequence to produce inaccurate performance.

Lee et al. [48] studied strategies for attacking the action space in MDP of reinforcement learning. They proposed two novel attack strategies, both of which were optimization-based approaches. The first is a Myopic Action Space (MAS) attack, which creates perturbations in a greedy manner without future considerations and then distributes these attacks across the action space dimensions. The second is a Look-ahead Action Space (LAS) attack, in which the attacker can make predictions, selects a designed sequence of future perturbations, and distributes the attacks across the action and temporal dimensions. The results show that LAS attacks have a greater influence on the agent than MAS attacks. Projected gradient descent (PGD) is used for the two formed optimal problems.

The approaches outlined above focus on attacking action space. Conversely, producing robust actions for reinforcement learning is also important and merits research attention.

Chen et al. [49] investigated reinforcement learning under action attack and considered actions to be robust in two situations. One is the Probabilistic Action Robust MDP (PR-MDP), with an adversary that adds a perturbation to the selected action; the other is the Noisy Action Robust MDP (NR-MDP),



in which an agent has a probability of taking an alternative adversarial action. PR-MDP and the NR-MDP can both be viewed as a zero-sum game. Subsequently, Policy Iteration (PI) schemes are used to solve these problems in order to reach Nash Equilibrium or convergence.

### 4.1.2 State Security in MDP

The state space, which is another important part of MDP, has a similar character to the action space in reinforcement learning systems. Both of them can affect the evaluation of policy and the environment. Methods of perturbing the state space or deceiving agents in this space have also been investigated, with particular attention paid to the dynamics of the agent cases. There are two common approaches to crafting state-space perturbations of reinforcement learning: model-based means and optimization based-means.

Zhao et al. [40] studied adversarial sample attacks in reinforcement learning using a model-based mean. They attempted to formulate an approximate sequence-to-sequence (seq2seq) model to predict a single agent action or a sequence of future agent actions, and produce adversarial samples to affect the state in MDP. They considered a full black-box attack situation, in which the attacker has no message of the agents, regardless of the training parameters or training methods of the agents. Finally, the adversarial samples are used to trigger a trained agent to misbehave at a specific time. This is a new concept in adversarial sample reinforcement learning that involves a time-bomb attack.

For certain high-dimensional continuous state space situations, computing an optimization-based approach may be a more suitable means of generating attacks for a target agent compared to training another model.

Garrett et al. [50] proposed an optimization method to formulate an attack that can directly influence the state(s). These attacks can affect the learning policy, cause the agent to act sub-optimally, and destabilize the cyber-physical system; this is made possible by the assumption that, in some cases, a foe can manipulate sensing/actuation signals. These authors used a Z table to conduct an attack. The Z table is a measure of how effective an adversarial effect to a state is versus the cost to perform the effect. Then using iteration to get the target state(s).

Models also can be used as an auxiliary means of making predictions that improve the attack methods. Sun et al. [35] proposed adversarial attacks that add perturbation into the agent's observation state. The goal of adversarial attacks against Deep Reinforcement Learning (DRL) is to inject a small set of adversarial samples in critical moments. For simplicity, this problem can be regarded as two sub-problems: that of "when to attack", and that of "how to attack". Critical point attack and antagonist attack are two models proposed to predict environment states and accordingly discover the critical steps and locations for perturbations. Existing adversarial example techniques are then used to compute and add perturbations in the selected critical moments.

Ye et al. [51] proposed a model-based self-advising method for multi-agent learning. This method enables the agents with the same ability when asked

for advice on an unfamiliar state. The idea is to train a model based on states similar to a certain state. These authors provided a defined conception of "Similar States" and adopted a deep neural network to train a teacher agent on states both unfamiliar and familiar to the student. This method produces an improvement in learning and more robust performance with a much lower communication overhead.

### 4.1.3 State and Action Security in MDP

Sometimes, action and state space can be considered at the same time, at which point they are regarded as state-action pairs. These pairs can affect the reward function and potentially skew the policy so that the agent makes poor decisions in the Markov process.

Dai et al. [52] proposed a reinforcement learning method that considers safe exploration by evaluating the risk level of each state-action pair, then recording the most dangerous state-action pairs based on the security performance metrics in a blacklist. This algorithm employs a modified Boltzmann distribution based on the Q-values and the risk levels to choose an action. Moreover, a convolutional neural network was adopted to weigh the long-term risk levels of each state-action pair, while transfer learning was chosen to reduce the initial explorations in initial parameters learning.

We can observe that almost all the papers mentioned above consider adding perturbations to the target to influence a system. Moreover, the popular method of identifying a good perturbation is to regard the issue as an optimization problem. As a result, optimization-based approaches such as the gradient method and some intelligent approaches such as imitation learning have been selected for these purposes. Optimization-based approaches are traditional algorithms used to solve optimization problems, which are very efficient and intuitive; however, these methods may be useless for complex problems. In contrast, intelligent approaches can deal with these problems, thus these approaches can extend to more areas. Nevertheless, intelligent approaches may require more data for learning and more computing resources.

## 4.2 Security of Environment

In reinforcement learning systems, an agent interacts with the environment and then adjusts its learning results based on the information obtained from the environment. The environment includes the transition dynamics in MDP and surrounding environment situations; examples of the latter include the surrounding architecture or road conditions encountered by the agent in the real world. Affecting the transition dynamics in MDP and the surrounding environment situations can also influence the reward or state of environment input received by the agent; thus, it can mislead the agent into taking incorrect actions, and can also affect the policy.

**Table 3** Comparison of adding perturbations for state and action security

Paper	Target	Impact	Key Idea	Representative Methods	Tradeoffs and Shortcomings
Lee et al. [48]	Action	Reward	Algebra-based optimization	Optimization-based approaches Projected gradient descent	May be useless in non-convex problems
Garrett et al. [50]	State	System destabilization	approaches	Z Tables	
Zhao et al. [40]	State	Policy Action	Intelligent approaches	Imitation learning	May require more computing resources
Sun et al. [35]	State	Reward, Action		Prediction model Neural network	

Rakhsha et al. [38] and Chan et al. [53] focus on studying attacks on the transition dynamics and the features obtained from the environment respectively. Wang et al. [22], Li et al. [54] and Zhai et al. [55] focused on policy robustness in disturbed environmental conditions. All three of the above ideas are for single-agent areas. However, Li et al. [41] and Lin et al. [39] focused on the environment issue in multi-agent fields. Li et al. [41] aimed at building robust policies under a non-stationary environment, while Lin et al. [39] proposed attacks on the feature from the environment that mislead the agent into taking a target action.

#### 4.2.1 Environment security of a single agent

The issue of an attacker who harms or poisons the learning environment was researched by Rakhsha et al. [38]. These authors focused on data poisoning attacks that can manipulate the rewards or the transition dynamics in the reinforcement learning (RL) algorithm, based on the understanding that reinforcement learning agents aim to maximize their average reward in undiscounted infinite-horizon settings. To fabricate the attacks, optimal methods with constraints and bounds were used. This paper considered online learning settings with poisoned feedback, in which the agent uses a regret-minimization mechanism to learn a policy while considering the different attack costs for online learning settings. In addition, it also studied offline settings in which the agent is formulating plans in a poisoned environment.

In many cases, the surrounding environment information is complicated or high-dimensional, meaning that pre-management is required. Thus, the opponent can not only focus on the surrounding environment data directly, but can also aim at the data after some calculating steps (for example, extracting features).

Chan et al. [53] studied the adversarial attack strategy against DRL by crafting an adversarial sample that perturbs the features to efficiently affect the cumulative reward. A static reward impact map is presented to measure the influence on the cumulative reward made by inputting features, that have slight changes. Then, using the "reward impact map" measures the importance of a feature; subsequently, select suitable actions. Finally, a crafted adversarial sample based on the gradient function attacks the sample by perturbing the features.

The studies above focus on attacks. However, it is also important to develop a more reliable and stable reinforcement learning system for the surrounding situation. Accordingly, some approaches were developed and tested in a perturbed or poisoned environment.

Wang et al. [22] developed a mechanism that considers reinforcement learning with safety falsification methods. This framework is a falsification-based robust adversarial reinforcement learning (FRARL) framework that trains the policy in the new adversarial environment; thus, the system can perform as an adversarial reinforcement learning mechanism to enhance the robustness of trained policies. It is the first generic mechanism that combines temporal

logic falsification with adversarial learning to improve policy robustness. A cross-entropy method is used to get the initial conditions and input sequences.

Li et al. [54] considered policy learning for Robust Markov Decision Processes (RMDP) from another unique perspective. Specifically, rather than concentrating on attack problems, these authors focused on the robustness of the environment to simulator domain mismatch in real application scenarios. They treated the mismatch as a perturbation and established the goal of finding a robust policy that ensures a near-optimal reward against the worst-case perturbation. A two-player zero-sum game was developed that considers the perturbation as an adversarial player, and Nash Equilibrium (NE) was used to find the robust policy.

Zhai et al. [55] also studied the problem of the differences between simulated and real environment, which may reduce the performance of the learned policies. These authors also modeled environmental differences as adversarial disturbances and constructed a two-player-zero-sum game between the normal and adversarial agents. However, such a method may increase the difficulty of the training domain. Consequently, this paper also considered certain constraints in the adversarial architecture and used a data-driven Lyapunov network to ensure the stability of the system during training.

We can observe that single-agent reinforcement learning is delicate and sensitive in the training surrounding situation. Notably, this issue is even more pronounced in multi-agent scenarios. Researchers have accordingly studied these problems to make the learned models more robust for multi-agents.

#### **4.2.2 Environment security of multi-agents**

Li et al. [41] studied the problem of training robust deep reinforcement learning agents with continuous actions in the multi-agent learning setting. They proposed a Minimax Multi-Agent Deep Deterministic Policy Gradient algorithm (M3DDPG) to improve the robustness of the multi-agent reinforcement learning system. Their approach introduced the Minimax Optimization idea to update policies considering the worst situation. When this approach is applied, each agent operates under the assumption that all other agents are acting adversarially. The algorithm was based on Multi-Agent Deep Deterministic Policy Gradient (MADDPG), a decentralized policy and a centralized critic framework, used to obtain a Q function. These authors further proposed Multi-Agent Adversarial Learning (MAAL), which can approximate the non-linear Q function using a locally linear function to solve this optimization problem.

Lin et al. [39] studied attacks on cooperative multi-agent reinforcement learning (c-MARL), which is a necessary element of improving the robustness of the algorithm. In this study, perturbing agents' observations with an adversarial example can mislead the agents and minimize the value of a team's total reward. This special attack method comprises two steps. First, the adversary selects actions that can minimize the total team reward using a policy network

trained with reinforcement learning. Next, the rival perturbs the agent's observation by using targeted adversarial examples and gradient-based methods to make the agent take specific actions.

Some of the papers discussed above focus on attacks, while others are about defending the environment. Attacks of this kind almost always utilize adversarial samples, while defenses against such attacks are always robust adversarial learning methods that consider training the agent in a worse situation to obtain a robust model.

### 4.3 Security of reward function in MDP

In many reinforcement learning (RL) applications, the agent extracts reward signals from the user or environment. A reward is crucial for a Markov decision process, as it indicates the feedback received by agents when they take certain actions in certain states. Ultimately, the reward is used to decide which action is optimal. Thus, adversaries tend to attack the reward functions of MDP in reinforcement learning.

#### 4.3.1 Attack reward function in MDP

The reward-poisoning attacks issued against reinforcement-learning agents were investigated by Zhang et al. [56]. In this paper, these authors studied the training-time reward poisoning attack problem, which involves crafting environmental rewards and forces the reinforcement learning agent to learn a nefarious policy. They regarded the reward shaping task as an optimal control problem on a higher-level attack MDP. They also characterized conditions in situations where such attacks are useless, as well as provided upper bounds on the attack cost when an attack is feasible.

The concept discussed above involves learning to attack the reinforcement learning systems through reward. A robust reward signal therefore becomes important, and has accordingly attracted significant attention from researchers.

#### 4.3.2 Protecting reward function in MDP

Li et al. [57] focused on the problems of reward signal sparsity and instability in the field of dialogue generation. Dialogue reward learning with adversarial inverse reinforcement learning (DG-AIRL) is proposed to address this issue. This mechanism is a sequence-to-sequence (Seq2Seq) model that adopts adversarial imitation learning to enable the model to give human-like dialogue responses, and further designs a specific reward function structure to measure the reward of each word in the generated sentences. An entropy regularization term is also used to improve training stability.

The above-mentioned paper aimed at attacking the reward function using a poisoning attack that selected an optimal control method. It is possible that learning methods could be used to handle such a problem. Moreover, the adversarial sample which is another popular attack algorithm, can also be

**Table 4** Comparison attacks and defences for environment security

Strategies	Paper	Categories	Impact	Representative Methods
Adversarial sample	Chan et al. [53]	Attack	Cumulative reward	Sliding-window method Gradient function
Robust adversarial learning	Lin et al. [39]	Attack	Action	Gradient-based methods
	Wang et al. [22]	Defense	Robust policy	Cross-entropy method Actor-critic architecture
	Li et al. [41]	Defense	Robust policies	Minimax optimization End-to-end learning approach

applied to reward function attacks in future work. In addition, the adversarial inverse reinforcement learning method can successfully improve the quality of intelligence in reward inferring, and may be extended in other areas.

A summary of research into security in reinforcement learning is presented in Table 5.

## 5 Privacy in reinforcement learning

In this section, we will discuss privacy in reinforcement learning problems such as obtaining or inferring private information and protecting user privacy, with a particular focus on the MDP perspective. This section can be divided into three subsections. The first explores the privacy of state and action in MDP. The second one addresses the privacy of the environment, which encompasses the transition dynamics in MDP and the surrounding environment situations. The last part examines the privacy of the reward function in MDP.

### 5.1 Privacy of state and action in MDP

The agent and the environment interact through states and actions, which may be stolen by adversaries and cause privacy leakage problems.

States and actions in MDP encompass the state of the environment, the action of the agent, and other concepts based on state and action. Vietri et al. [58] aim at protecting the state and action directly in MDP. Wang et al. [37], Sakuma et al. [23], Ye et al. [59], Cheng et al. [60], and Chowdhury and Zhou [61] focused on protecting the Q-functions based on states and actions.

#### 5.1.1 Using differential privacy to ensure the privacy of state and action in MDP

For privacy protection problems, differential privacy is a popular model that has been widely used in many areas. Differential privacy considers the worst-case situation in which attackers know all the data except for a new record in a dataset. It can ensure that any individual record being stored in or removed from a dataset makes little difference to the dataset's analytical output [62].

The work of Wang et al. [37] considered the use of differential private methods for reinforcement learning in continuous spaces, with a focus on protecting the value function approximator. The algorithm added Gaussian process noise to the corresponding action-state value function of deep Q-learning, which can satisfy differential privacy guarantees at every iteration. It also chose the reproducing kernel Hilbert space (RKHS) embedding common neural networks for the nonlinear value function.

Ye et al. [59] also adopted differential privacy to study issues in reinforcement learning. These authors proposed a novel differential advising in multi-agent reinforcement learning inspired by the differential privacy mechanism. Using this approach, an agent can take advice produced with reference to



Table 5 Summary of research addressing security in reinforcement learning

Subsection	Papers	Target	Impact	Strategies	Representative Methods
Security of state and action in MDP	Lee et al. [48]	Action	Reward	Perturbations	Optimization-based approaches Projected gradient descent
	Chen et al. [49]	Action	Policy	Action robustness	Zero-sum game
	Zhao et al. [40]	State	Policy Action	Perturbations	Nash equilibrium
	Garrett et al. [50]	State	System destabilization Reward, Action	Perturbations	Imitation learning
	Sun et al. [35]	State	Action	Perturbations	Z tables
	Ye et al. [51]	State	Action	Model learning	Prediction model Neural network
	Dai et al. [52]	State-action	Policy	Safe exploration	Deep neural network Convolutional neural network
					Transfer learning
Security of environment in MDP	Rakhsha et al. [38]	Transition dynamics and/or rewards	Policy	Data poisoning	Optimization problems having constraints
	Chan et al. [53]	Features	Cumulative reward	Adversarial sample	Sliding-window method Gradient function
	Wang et al. [22]	Environment conditions	Robust policy	Robust adversarial learning	Cross-entropy method Actor-critic architecture
	Li et al. [41]	Non-stationary environment	Robust policies	Robust adversarial learning	Minimax optimization
	Lin et al. [39]	Features	Action	Adversarial sample	End-to-end learning approach Gradient-based methods
	Li et al. [54]	Environment	Policy	Two-player zero-sum game	Nash equilibrium
	Zhai et al. [55]	Environment	Policy	Two-player zero-sum game	Nash equilibrium Lyapunov network
Security of reward function in MDP	Zhang et al. [56]	Reward	Policy	Poisoning attack	Optimal control problems
	Li et al. [57]	Reward	Policy	Adversarial inverse reinforcement learning	Imitation learning Entropy regularization term

a slightly different state. It also added Laplace noise to the agent advice  $Q(s)$ , which can provide the agent with more reliable data for use in decision-making.

The knowledge transfer problem in multi-agent reinforcement learning was studied by Cheng et al. [60]. They proposed a Differential knowledge Transfer with relevance Weight (DTW) algorithm, which also chose differential privacy to handle this problem. DTW was embedded into the multi-agent reinforcement learning algorithm to add differential noise and relevance weights to the  $Q$ -value. This model can expand the knowledge set and reduce the influence of negative transfer.

Far less attention has been paid to addressing privacy in reinforcement learning problems, compared with private bandit algorithms. Vietri et al. [58] proposed the first reinforcement learning algorithm for regret minimization with the JDP guarantee. These authors designed an algorithm named the Private Upper Confidence Bound algorithm (PUCB) for reinforcement learning, which used differentially private guaranteeing of the Laplace mechanism to protect the information of training data, and moreover imposed lower bounds on the regret and a smaller number of sub-optimal episodes. It computed the policy of the reinforcement learning algorithm using private counts such as  $\hat{n}(s, a, h)$  and  $\hat{m}(s, a, s', h)$ , which are the number of times the agent has taken action  $a$  in state  $s$  at time  $h$  and the number of times the agent has taken action  $a$  in state  $s$  at time  $h$  and transitioned to  $s'$  respectively. This method satisfies the joint differential privacy (JDP) with lower-bounds sample complexity and regret of probably approximately correct (PAC).

Chowdhury and Zhou [61] also considered the regret bounds of reinforcement learning. They proposed two general frameworks for designing private, optimistic reinforcement learning algorithms, one for policy optimization and another for value iteration, that also satisfied JDP and LDP requirements. They designed the counts returned by the privatizer, which depend on users' states and actions to calculate the private mean empirical costs and private empirical transitions. Examples include the count  $N_h^k(s, a)$  and its privatized versions  $\tilde{N}_h^k(s, a)$ , which denote the number of times that the agent has visited state-action pair  $(s, a)$  at step  $h$  before episode  $k$  and are similar to the counts in [58]. The frameworks also can obtain sublinear regret guarantees.

### 5.1.2 Using cryptography to improve the privacy of state and action in MDP

Sakuma et al. [23] studied the issue of privacy in distributed reinforcement learning (DRL), devising privacy-preserving reinforcement learning algorithms using an additive homomorphic cryptosystem. The  $Q$ -values are encrypted, allowing the addition of encrypted values without requiring their decryption. Data partitioning by time and by observation were considered, and random action selection was used for these two aspects.

Differential privacy and cryptography are both prevalent privacy models, each with its advantages and disadvantages. Differential privacy may be more flexible and can be used in many fields due to its ability to achieve a balance

between data utility and privacy. However, this method may sometimes be unreliable, leading to privacy leakage and the sacrifice of some data utility. In contrast, cryptographic techniques are very reliable, but at the same time, they require a great deal of calculation and may reduce the efficiency of data sharing.

## 5.2 Privacy of environment

Interaction and data sharing occur frequently between the agent and the environment, and the environment is closely connected with other information in MDP. Hence, there is also a high probability of environmental privacy leakage. Opponents will focus on stealing the transition dynamics data of the environment or information of the surrounding conditions (The surrounding conditions denote the area in which the agents are trained and/or to which they are applied).

The privacy leakage problem in deep reinforcement learning was studied by Pan et al. [18]. These authors focused on the problem of environment dynamics search with the goal of inferring the environment. Genetic algorithm and shadow policies were selected for optimal policy selection and candidate inference respectively. This private environment information leaking problem was considered under two different scenarios. In the first, the attacker has no knowledge about the training surrounding situations, and the environments just with common constraints. In the second, the attacker has access to a set of potential candidates of the training environment dynamics.

Ye et al. [63] studied private information leakage of Multi-Agent planning for logistic-like problems. These authors adopted a planning approach that employed a reinforcement learning algorithm to make a plan to find the optimal route from the initial state to the goal state. This paper, proposed an approach adopting the differential privacy technique to achieve strong privacy preservation in multi-agent environments. It also used the concept of a “privacy budget ” to control the communication overhead.

Zhou et al. [64] proposed to protect users’ sensitive and private information by considering regret minimization in large state and action spaces. Their work used the notion of joint differential privacy (JDP) and considered MDPs by means of linear function approximation. It further proposed two algorithms that applied the Gaussian mechanism to the information of the environment to protect the features. The two proposed privacy reinforcement learning algorithms are based on value iteration and policy optimization.

## 5.3 Privacy of reward function in MDP

Reinforcement learning is a framework within which an agent learns a behaviour policy through interacting with the environment and responding to positive and negative rewards [65]. The reward function always determines the amount of the reward and when it is given. Thus, the reward function is a

**Table 6** Comparison of methods for defending the privacy of state and action in MDP

Key idea	Papers	Target	Representative Methods	Pros	Cons
Cryptographic solutions	Sakuma et al. [23]	Q-values	Additive homomorphic cryptosystem	Reliable	High resource consumption Low data sharing efficiency
	Wang et al. [37]	Q-values	Gaussian process noise Common neural networks		
Differential privacy	Ye et al. [59]	Q-values	Laplace mechanism	Efficient	Unsafe in some cases
	Cheng et al. [60]	Q-values	Differentially exponential noise Relevance weights		
	Vietri et al. [58]	States, actions	Joint differential privacy Optimistic strategy		
	Chowdhury and Zhou [61]	States, actions	Joint differential privacy Local differential privacy		

key element of giving the reward, and is accordingly very important for reinforcement learning. However, it is likely that observers can attempt to infer information about the policy from the reward function.

Liu et al. [66] aimed to preserve the privacy of a reward function in reinforcement learning. Specifically, these authors proposed two methods with more general dissimulation models to preserve reward privacy: the Ambiguity Model, in which the agent selects actions that maximize the entropy based on ambiguity, and the intention recognition model, which takes action selection as a weighted sum of honest and ‘irrational’ behaviour. These methods both use pre-trained Q-functions and produce a policy that makes it hard to use inverse reinforcement learning or imitation learning. Thus, an observer is difficult to obtain the reward function.

Ye et al. [24] addressed issues with Multi-agent Advising Learning and applied reinforcement learning to deal with the packet routing problem. These authors adopted differential privacy to reduce the impact of the malicious agent by adding Laplace noise to the accumulated reward to protect the information of each agent. The privacy budget concept was then used to control the communication overhead, which can improve the learning performance.

Fu et al. [67] proposed an adversarial inverse reinforcement learning algorithm to acquire a robust reward for changes in dynamics. Inverse reinforcement learning focuses on the problem of inferring an expert’s reward function from demonstrations, and this paper combines such a mechanism with adversarial learning, which can improve the robustness of the algorithm. This approach enables the proposed algorithm to learn policies even in environments that undergo great changes during training; thus, it achieves better performance than prior IRL methods in continuous, high-dimensional situations with unknown dynamics.

As differential privacy is always adopted to establish a mathematical way of guaranteeing data privacy in reinforcement learning, and considering that inverse reinforcement learning is applied to inferring the reward function from demonstrations and providing rewards to the learning system, Prakash et al. [68] investigated the existing set of privacy techniques for reinforcement learning and proposed a new Privacy-Aware Inverse reinforcement Learning (PRIL) analysis framework, which is a new form of privacy attack that targets the private reward function. This reward attack attempts to reconstruct the original reward from a privacy-preserving policy (such as differential privacy) using an inverse reinforcement algorithm. The results showed that privacy in the policy domain does not translate to privacy in the reward domain, as the reconstruction error is independent of the  $\epsilon$ -DP budget.

Privacy problems related to the environment information and reward function are shown above. Environment components, the surrounding situation and the reward function are considered. We can observe that differential privacy is the most popular method, and moreover that some intelligent algorithms have been developed in the privacy arena. Intelligent algorithms can handle more complex problems and situations, even if little information about the model

and the goal is available; thus, they can expand the horizons of privacy protection. We can further observe that inverse reinforcement learning has been adopted to address these privacy problems [64, 67]. In fact, however, inverse reinforcement learning is more frequently used to approach the problem of inferring an expert’s reward function from demonstrations and to provide the reward to the learning system rather than to tackle security and/or privacy issues. However, the concept of ”inferring a reward function” naturally prompts thoughts of privacy leakage and unreliable models. Consequently, it might be possible to consider the security and privacy problems of inverse reinforcement learning, like the work in [68].

Table 8 provides a concise summary of the research discussed above.

## 6 Future works

This paper focuses on security and privacy in reinforcement learning. Developing a stable and reliable algorithm is an important direction of reinforcement learning. Attacks and defenses of MDP in this area are both key elements. There are numerous related research avenues that could be pursued in future.

**Attacking various components of MDP** In the papers discussed above, opponents may attack only one or two elements of MDP. We may study attacking many different elements in MDP in future, along with the impact caused by more elements. We could also investigate the combined action of attacking multiple elements simultaneously; it is possible that the combined effect will exceed the total impact of separate attacks.

In addition, we can further study the relationship between the elements with the hope of more strongly influencing the system performance. We can infer other information about the system elements based on some elements in MDP, or find the most efficient attack point or moment based on the relationship between these elements.

Furthermore, we could also discuss issues related to the limited knowledge possessed by attackers. For example, if an attacker possesses partial state information and partial Dynamic Transition probability information, it may be possible to produce attacks based on this limited knowledge alone and then combine the analysis of the relationship between these two elements.

**Improving robustness of reinforcement learning** In real situations, the environment is complex and continuously changing. We can consider fault-tolerant control to develop a more robust algorithm that is capable of handling this complex and changing environment. For example, we could consider the situation in which the agent receives incorrect information, or fails to receive some part of the required information, because of issues with the sensors. We could also consider the situation in which the connection quality is poor, preventing the agent from receiving information in a timely fashion and consequently causing it to make sub-optimal decisions.

An adaptive fault-tolerant control (FTC) approach for MIMO nonlinear discrete-time systems was proposed by Liu et al. [69]. In this paper, abrupt

**Table 7** Comparison of methods for defending environment and reward function privacy in MDP

Key idea	Papers	Impact	Strategies	Representative Methods	Tradeoffs and Shortcomes
Intelligent means	Liu et al. [66]	Q-values	Reinforcement learning	Ambiguity model Intention recognition	Higher computational consumption; can only defend against certain specific attacks
	Fu et al. [67]	Policy	Adversarial learning	Inverse reinforcement learning	
Differential privacy	Ye et al. [63]	Q-values Reward	Differential privacy	Laplace noise Privacy budget	Higher communication overhead
	Zhou et al. [64]	Policy	Joint differential privacy	Gaussian mechanism	
	Ye et al. [24]	Policy	Differential privacy	Laplace noise	

Table 8 Summary of research addressing the Privacy in reinforcement learning

Subsection	Papers	Target	Impact	Strategies	Representative Methods
Privacy of state and action in MDP	Sakuma et al. [23]	Q-values	States, actions, reward	Cryptographic solutions	Additive homomorphic cryptosystem
	Wang et al. [37]	Q-values	Reward	Differential privacy	Gaussian process noise Common neural networks
	Ye et al [59]	Q-values	Policy	Differential advising method	Laplace mechanism
	Cheng et al. [60]	Q-values	Policy	Differential transfer learning method	Differentially exponential noise Relevance weight
Privacy of environment in MDP	Vietri et al. [58]	States, actions	States, actions, reward	Differential privacy	Joint differential privacy Optimistic strategy
	Chowdhury and Zhou [61]	States, actions	Transition probabilities	Differential privacy	Joint differential privacy Local differential privacy
	Pan et al. [18]	Transition dynamics	Surrounding environment situations	Environment dynamics search	Genetic algorithm Shadow policies
	Ye et al. [63]	Environment	Q-values, reward	Differential privacy	Laplace noise Privacy budget
Privacy of reward function in MDP	Zhou et al. [64]	Environment	Policy	Joint differential privacy	Gaussian mechanism
	Liu et al. [66]	Reward function	Q-values	Reinforcement learning	Ambiguity model Intention recognition
	Ye et al. [24]	Reward function	Policy	Differential privacy	Laplace noise
	Fu et al. [67]	Reward	Policy	Adversarial learning	Inverse reinforcement learning



faults and incipient faults are both taken into account. Li et al. [70] also focused on the adaptive fault-tolerant tracking control problem, and further considered the influence of the dead zones and actuator faults on the control performance. Future work could thus take specific faults into account based on the actual situation.

**Security and privacy issues of smart methods in reinforcement learning** Many intelligent methods have been used for reinforcement learning problems. These methods may be used in complex problems in which it is difficult to obtain the goal or the model.

Inverse reinforcement learning (IRL) is one such intelligent method that is used to infer the reward function of the reinforcement learning model. Its application fields include video games, in which it is often more difficult to design a reward function that describes the behaviors and yields an optimal policy than to provide demonstrations of the target behaviour in the video games. Tucker et al. [71] used inverse reinforcement learning algorithms to infer a reward from demonstrations; this approach utilized CNNs to deal with high-dimensional video games. Neu et al. [72] investigated the application of IRL algorithms to parser training problems, and were able to automatically find a reward function that matched the training set. IRL may be a good method for use in building a better reinforcement learning model, as it enables rewards to be obtained automatically. However, it also may be used to attack a system to infer the reward function. Moreover, to obtain rewards, IRL also needs data to train; as a result, privacy leakage problems may arise [68].

In addition, while such intelligent algorithms can improve the handling of certain problems, it might be preferable to focus on specific convergence analysis rather than simply learning in future work (for example, studying the regret bounds of reinforcement learning [61]).

## 7 Conclusion

In this paper, we investigated security and privacy in reinforcement learning. We analyzed the targets and impacts from the perspective of MDP, and review existing research based on the elements of MDP. Specifically, we recognize attacks and defenses of state and action, environment and reward.

We also described the strategies and representative methods of security and privacy issues, facilitating a clear understanding of what method is used for which object of the MDP. We conducted an analysis combining those methods with the character of the elements of the MDP tuple. We went on to discuss the advantages and limitations of the studies, along with potential future directions of research into security and/or privacy in reinforcement learning.

In summarizing the recent research into security and privacy in reinforcement learning, the following important findings can be extracted:

- Every element in MDP can be attacked, and can thus affect the overall process.

- Security and privacy issues in reinforcement learning are always regarded as optimal problems, and mathematical optimization methods and some intelligent learning algorithms are adopted to deal with the optimal problems.

Based on these findings, we suggest three directions for future research:

- Considering the elements in MDP. Investigating different components of MDP and the relationships between them may help to alleviate the security and privacy problems in reinforcement learning.
- Considering the real environment and actual situations of the learning system is an interesting direction; for example, exploring equipment faults and environmental disturbances.
- Given the many intelligent methods used in the various areas of reinforcement learning research, it would seem wise to discuss the potential problems that may occur as a result, and to improve the performance of those intelligent methods.

## References

- [1] Ying, Z., Zhang, Y., Cao, S., Xu, S., Liu, X.: Oidpr: Optimized insulin dosage based on privacy-preserving reinforcement learning. In: 2020 IFIP Networking Conference (Networking), pp. 655–657 (2020). IEEE
- [2] Meng, T.L., Khushi, M.: Reinforcement learning in financial markets. *Data* **4**(3), 110 (2019)
- [3] Kober, J., Bagnell, J.A., Peters, J.: Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* **32**(11), 1238–1274 (2013)
- [4] Gosavi, A.: Reinforcement learning: A tutorial survey and recent advances. *INFORMS Journal on Computing* **21**(2), 178–192 (2009)
- [5] Lei, L., Tan, Y., Zheng, K., Liu, S., Zhang, K., Shen, X.: Deep reinforcement learning for autonomous internet of things: Model, applications and challenges. *IEEE Communications Surveys & Tutorials* **22**(3), 1722–1760 (2020)
- [6] Li, L., Chu, W., Langford, J., Schapire, R.E.: A contextual-bandit approach to personalized news article recommendation. In: *Proceedings of the 19th International Conference on World Wide Web*, pp. 661–670 (2010)
- [7] Yu, C., Liu, J., Nemati, S., Yin, G.: Reinforcement learning in healthcare: A survey. *ACM Computing Surveys (CSUR)* **55**(1), 1–36 (2021)
- [8] Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research* **17**(1),

1334–1373 (2016)

- [9] Deng, Y., Bao, F., Kong, Y., Ren, Z., Dai, Q.: Deep direct reinforcement learning for financial signal representation and trading. *IEEE transactions on neural networks and learning systems* **28**(3), 653–664 (2016)
- [10] Pan, X., You, Y., Wang, Z., Lu, C.: Virtual to real reinforcement learning for autonomous driving. *arXiv preprint arXiv:1704.03952* (2017)
- [11] François-Lavet, V.: Contributions to deep reinforcement learning and its applications in smartgrids. PhD thesis, Universite de Liege, Liege, Belgique (2017)
- [12] Uprety, A., Rawat, D.B.: Reinforcement learning for iot security: A comprehensive survey. *IEEE Internet of Things Journal* **8**(11), 8693–8706 (2020)
- [13] Bellifemine, F.L., Caire, G., Greenwood, D.: Developing Multi-agent Systems with JADE. John Wiley & Sons, ??? (2007)
- [14] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., *et al.*: Human-level control through deep reinforcement learning. *nature* **518**(7540), 529–533 (2015)
- [15] Zhu, T., Li, G., Zhou, W., Philip, S.Y.: Differentially private data publishing and analysis: A survey. *IEEE Transactions on Knowledge and Data Engineering* **29**(8), 1619–1638 (2017)
- [16] Park, J., Kim, D.S., Lim, H.: Privacy-preserving reinforcement learning using homomorphic encryption in cloud computing infrastructures. *IEEE Access* **8**, 203564–203579 (2020)
- [17] Xiao, Z., Xiao, Y.: Security and privacy in cloud computing. *IEEE communications surveys & tutorials* **15**(2), 843–859 (2012)
- [18] Pan, X., Wang, W., Zhang, X., Li, B., Yi, J., Song, D.: How you act tells a lot: Privacy-leaking attack on deep reinforcement learning. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 368–376 (2019)
- [19] Rakhsha, A., Radanovic, G., Devidze, R., Zhu, X., Singla, A.: Policy teaching in reinforcement learning via environment poisoning attacks. *Journal of Machine Learning Research* **22**(210), 1–45 (2021)
- [20] Huang, Y., Zhu, Q.: Deceptive reinforcement learning under adversarial manipulations on cost signals. In: *International Conference on Decision*

- and Game Theory for Security, pp. 217–237 (2019). Springer
- [21] Behzadan, V., Munir, A.: Vulnerability of deep reinforcement learning to policy induction attacks. In: International Conference on Machine Learning and Data Mining in Pattern Recognition, pp. 262–275 (2017). Springer
- [22] Wang, X., Nair, S., Althoff, M.: Falsification-based robust adversarial reinforcement learning. In: 2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA), pp. 205–212 (2020). IEEE
- [23] Sakuma, J., Kobayashi, S., Wright, R.N.: Privacy-preserving reinforcement learning. In: Proceedings of the 25th International Conference on Machine Learning, pp. 864–871 (2008)
- [24] Ye, D., Zhu, T., Zhou, W., Philip, S.Y.: Differentially private malicious agent avoidance in multiagent advising learning. *IEEE transactions on cybernetics* **50**(10), 4214–4227 (2019)
- [25] Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT press, ??? (2018)
- [26] Yu, S., Chen, X., Zhou, Z., Gong, X., Wu, D.: When deep reinforcement learning meets federated learning: Intelligent multitimescale resource management for multiaccess edge computing in 5g ultradense network. *IEEE Internet of Things Journal* **8**(4), 2238–2251 (2020)
- [27] Chen, T., Liu, J., Xiang, Y., Niu, W., Tong, E., Han, Z.: Adversarial attack and defense in reinforcement learning-from ai security view. *Cybersecurity* **2**(1), 1–22 (2019)
- [28] Wu, Y., Wang, Z., Ma, Y., Leung, V.C.: Deep reinforcement learning for blockchain in industrial iot: A survey. *Computer Networks* **191**, 108004 (2021)
- [29] Chen, W., Qiu, X., Cai, T., Dai, H.-N., Zheng, Z., Zhang, Y.: Deep reinforcement learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials* (2021)
- [30] Bellman, R.: Dynamic programming princeton university press princeton. New Jersey Google Scholar (1957)
- [31] Littman, M.L., Dean, T.L., Kaelbling, L.P.: On the complexity of solving markov decision problems. arXiv preprint arXiv:1302.4971 (2013)
- [32] Barto, A.G., Sutton, R.S., Anderson, C.W.: Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on*

- systems, man, and cybernetics (5), 834–846 (1983)
- [33] Watkins, C.J.C.H.: Learning from delayed rewards (1989)
- [34] Watkins, C.J., Dayan, P.: Q-learning. *Machine learning* **8**(3-4), 279–292 (1992)
- [35] Sun, J., Zhang, T., Xie, X., Ma, L., Zheng, Y., Chen, K., Liu, Y.: Stealthy and efficient adversarial attacks against deep reinforcement learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 5883–5891 (2020)
- [36] Gandhi, D., Pinto, L., Gupta, A.: Learning to fly by crashing. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3948–3955 (2017). IEEE
- [37] Wang, B., Hegde, N.: Privacy-preserving q-learning with functional noise in continuous state spaces. *arXiv preprint arXiv:1901.10634* (2019)
- [38] Rakhsha, A., Radanovic, G., Devidze, R., Zhu, X., Singla, A.: Policy teaching via environment poisoning: Training-time adversarial attacks against reinforcement learning. In: *International Conference on Machine Learning*, pp. 7974–7984 (2020). PMLR
- [39] Lin, J., Dzeparoska, K., Zhang, S.Q., Leon-Garcia, A., Papernot, N.: On the robustness of cooperative multi-agent reinforcement learning. In: *2020 IEEE Security and Privacy Workshops (SPW)*, pp. 62–68 (2020). IEEE
- [40] Zhao, Y., Shumailov, I., Cui, H., Gao, X., Mullins, R., Anderson, R.: 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)*, pp. 16–24 (2020). IEEE
- [41] Li, S., Wu, Y., Cui, X., Dong, H., Fang, F., Russell, S.: Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 4213–4220 (2019)
- [42] Sehgal, A., La, H., Louis, S., Nguyen, H.: Deep reinforcement learning using genetic algorithm for parameter optimization. In: *2019 Third IEEE International Conference on Robotic Computing (IRC)*, pp. 596–601 (2019). IEEE
- [43] Arora, S., Doshi, P.: A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence* **297**, 103500 (2021)

- [44] Ye, D., Zhu, T., Shen, S., Zhou, W.: A differentially private game theoretic approach for deceiving cyber adversaries. *IEEE Transactions on Information Forensics and Security* **16**, 569–584 (2020)
- [45] Zhu, T., Ye, D., Wang, W., Zhou, W., Yu, P.S.: More than privacy: Applying differential privacy in key areas of artificial intelligence. *arXiv preprint arXiv:2008.01916* (2020)
- [46] Ye, D., Shen, S., Zhu, T., Liu, B., Zhou, W.: One parameter defense-defending against data inference attacks via differential privacy. *IEEE Transactions on Information Forensics and Security* (2022)
- [47] Alaya, B., Laouamer, L., Msilini, N.: Homomorphic encryption systems statement: Trends and challenges. *Computer Science Review* **36**, 100235 (2020)
- [48] Lee, X.Y., Ghadai, S., Tan, K.L., Hegde, C., Sarkar, S.: Spatiotemporally constrained action space attacks on deep reinforcement learning agents. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 4577–4584 (2020)
- [49] Tessler, C., Efroni, Y., Mannor, S.: Action robust reinforcement learning and applications in continuous control. In: *International Conference on Machine Learning*, pp. 6215–6224 (2019). PMLR
- [50] Garrett, I.Y., Gerdes, R.M.: Z table: Cost-optimized attack on reinforcement learning. In: *2019 First IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pp. 10–17 (2019). IEEE
- [51] Ye, D., Zhu, T., Zhu, C., Zhou, W., Philip, S.Y.: Model-based self-advising for multi-agent learning. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
- [52] Dai, C., Xiao, L., Wan, X., Chen, Y.: Reinforcement learning with safe exploration for network security. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3057–3061 (2019). IEEE
- [53] Chan, P.P., Wang, Y., Yeung, D.S.: Adversarial attack against deep reinforcement learning with static reward impact map. In: *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, pp. 334–343 (2020)
- [54] Li, J., Ren, T., Yan, D., Su, H., Zhu, J.: Policy learning for robust markov decision process with a mismatched generative mode. *arXiv preprint arXiv:2203.06587* (2022)

- [55] Zhai, P., Luo, J., Dong, Z., Zhang, L., Wang, S., Yang, D.: Robust adversarial reinforcement learning with dissipation inequation constraint (2022)
- [56] Zhang, X., Ma, Y., Singla, A., Zhu, X.: Adaptive reward-poisoning attacks against reinforcement learning. In: International Conference on Machine Learning, pp. 11225–11234 (2020). PMLR
- [57] Li, Z., Kiseleva, J., de Rijke, M.: Dialogue generation: From imitation learning to inverse reinforcement learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 6722–6729 (2019)
- [58] Vietri, G., Balle, B., Krishnamurthy, A., Wu, S.: Private reinforcement learning with pac and regret guarantees. In: International Conference on Machine Learning, pp. 9754–9764 (2020). PMLR
- [59] Ye, D., Zhu, T., Cheng, Z., Zhou, W., Philip, S.Y.: Differential advising in multiagent reinforcement learning. *IEEE Transactions on Cybernetics* (2020)
- [60] Cheng, Z., Ye, D., Zhu, T., Zhou, W., Yu, P.S., Zhu, C.: Multi-agent reinforcement learning via knowledge transfer with differentially private noise. *International Journal of Intelligent Systems* **37**(1), 799–828 (2022)
- [61] Chowdhury, S.R., Zhou, X.: Differentially private regret minimization in episodic markov decision processes. arXiv preprint arXiv:2112.10599 (2021)
- [62] Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A.: Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine* **34**(6), 26–38 (2017)
- [63] Ye, D., Zhu, T., Shen, S., Zhou, W., Yu, P.: Differentially private multi-agent planning for logistic-like problems. *IEEE Transactions on Dependable and Secure Computing* (2020)
- [64] Zhou, X.: Differentially private reinforcement learning with linear function approximation. arXiv preprint arXiv:2201.07052 (2022)
- [65] Sutton, R.S., Barto, A.G.: Reinforcement learning: an introduction mit press. Cambridge, MA **22447** (1998)
- [66] Liu, Z., Yang, Y., Miller, T., Masters, P.: Deceptive reinforcement learning for privacy-preserving planning. arXiv preprint arXiv:2102.03022 (2021)
- [67] Fu, J., Luo, K., Levine, S.: Learning robust rewards with adversarial inverse reinforcement learning. arXiv preprint arXiv:1710.11248 (2017)

- [68] Prakash, K., Husain, F., Paruchuri, P., Gujar, S.P.: How private is your rl policy? an inverse rl based analysis framework. arXiv preprint arXiv:2112.05495 (2021)
- [69] Liu, L., Wang, Z., Zhang, H.: Adaptive fault-tolerant tracking control for mimo discrete-time systems via reinforcement learning algorithm with less learning parameters. *IEEE Transactions on Automation Science and Engineering* **14**(1), 299–313 (2016)
- [70] Li, H., Wu, Y., Chen, M.: Adaptive fault-tolerant tracking control for discrete-time multiagent systems via reinforcement learning algorithm. *IEEE Transactions on Cybernetics* **51**(3), 1163–1174 (2020)
- [71] Tucker, A., Gleave, A., Russell, S.: Inverse reinforcement learning for video games. arXiv preprint arXiv:1810.10593 (2018)
- [72] Neu, G., Szepesvári, C.: Training parsers by inverse reinforcement learning. *Machine learning* **77**(2), 303–337 (2009)