GAIL-PT: A Generic Intelligent Penetration Testing Framework with Generative Adversarial Imitation Learning

Jinyin Chen[a,b], Shulong Hu[b], Haibin Zheng[b], Changyou Xing[c,*], Guomin Zhang[c]

[a]*Institute of Cyberspace Security, Zhejiang University of Technology, Hangzhou, China*
[b]*College of Information Engineering, Zhejiang University of Technology, Hangzhou, China*
[c]*College of Command and Control Engineering, Army Engineering University of PLA, Nanjing, China*

Abstract

Penetration testing (PT) is an efficient network testing and vulnerability mining tool by simulating a hacker's attack for valuable information applied in some areas, e.g., the host's operating and database systems. Compared with manual PT, intelligent PT has become a dominating mainstream due to less time-consuming and lower labor costs. Among the already proposed intelligent PT methods, reinforcement learning (RL) based PT has achieved state-of-the-art (SOTA) performance. Unfortunately, RL-based PT is still challenged in real exploitation scenarios because the agent's action space is usually high-dimensional discrete, thus leading to algorithm convergence difficulty. Besides, most PT methods still rely on the decisions of security experts. Addressing the challenges, for the first time, we introduce expert knowledge to guide the agent to make better decisions in RL-based PT, and propose a *Generative Adversarial Imitation Learning based generic intelligent Penetration Testing framework*, denoted as GAIL-PT, to solve the problems of higher labor costs due to the involvement of security experts and high-dimensional discrete action space. Specifically, first, we manually collect the state-action pairs to construct an expert knowledge base when the pre-trained RL / DRL model executes successful penetration testings. Second, we input the expert knowledge and the state-action pairs generated online by the different RL / DRL models into the discriminator of GAIL for training. At last, we apply the output reward of the discriminator to guide the agent to perform the action with a higher penetration success rate to improve PT's performance. Extensive experiments conducted on the real target host and simulated network scenarios show that GAIL-PT achieves SOTA penetration performance against DeepExploit in exploiting actual target Metasploitable2 and Q-learning in optimizing penetration path, not only in small-scale with or without honey-pot network environments, but also in the large-scale simulated network environment. The code of GAIL-PT is open-sourced at *https://github.com/Shulong98/GAIL-PT//*.

Keywords: Penetration testing, deep reinforcement learning, generative adversarial imitation learning, Metasploitable2, DeepExploit, Q-learning

*Corresponding author
Email addresses:* chenjinyin@zjut.edu.cn (Jinyin Chen), 2112003213@zjut.edu.cn (Shulong Hu), haibinzheng320@gmail.com (Haibin Zheng), changyouxing@126.com (Changyou Xing), zhanggmwn@163.com (Guomin Zhang)

1. Introduction

Penetration Testing (PT) is a process of simulating the attack behavior of malicious hackers, launching authorized attacks on computer systems and networks to discover any security vulnerabilities that can be exploited [1][2]. Numerous PT methods have been proposed for assessing the security of network systems, such as host operating systems [21], database systems [22], network equipment systems [13] [30] and etc.. According to different testing techniques, PT methods are roughly categorized as manual PT, automatic PT and intelligent PT.

Manual PT depends on applying software exploits like running computer code to trigger a system vulnerabilities and taking further actions such as payloads that after a successful exploit, according to the security audit experts[33]. Generally speaking, they take full advantage of some penetration tools (e.g., Nessus [49], Nexpose [37]) combined with expert operations to accomplish the PT process. However, those tools rely on the professional knowledge of security experts to make decisions, so they usually can not update attack strategies accordingly when the environment changes. Additionally, with the complexity of systems grows, the task of manually assessing security will become quite a burden of labor-consuming and time-consuming. Consequently, the emergence of automatic PT methods alleviate this problem to a certain extent.

Existing approaches of the automatic PT include those mapping the results of vulnerability scanners to the corresponding penetration tools and those describing the PT process as an attack graph to solve the problem of the penetration path planning. The former method is integrated by vulnerability scanning, penetration attacking, payloads choosing and other modules into a framework to realize PT automatically, e.g., Metasploit [23][38], Core Impact [15]. However, some operations still require manual participation in the critical penetration steps, such as selecting proper exploit modules and payloads. In other words, it still relies on the attack rules carefully set in advance by programmers and security experts [31]; thus, it cannot learn to deal with dynamic and uncertain penetration environments to develop new attack strategies. The latter method of automatic PT is mainly based on attack graphs [32], by showing the attack sequence and the effect that an attacker may launch. In the graph, the attacker could use the relationship of vulnerabilities and exploits that are already grasp to choose an attack path. Nevertheless, since the attack graphs demands the complete system knowledge, it is difficult to apply to realistic dynamic penetration environments. To address the problem, Artificial Intelligence (AI) based PT methods come into being.

According to the AI techniques adopted for PT, AI-based intelligent PT methods can be mainly grouped into PT based on traditional Reinforcement Learning (RL) and Deep Reinforcement Learning (DRL) [29].Specifically, in PT based on traditional RL, such as modeling the attack process as a Partially Observable Markov Decision Process (POMDP) [39][40], the attacker gradually observes and models the computer configuration as the attack progresses. It can be well-aimed at a single host in practice, but due to POMDP's high computational complexity, it is not scalable for large networks [43]. Besides, the

Q-learning algorithm is also applied to the Capture The Flag (CTF) competition [26] and the penetration path optimization [56]. However, these methods are also challenged in real-world scenarios since they are hard to involve in large-scale networks. Besides, the RL agent needs to comprehend the complete network topology in advance, which is a strong assumption in the real environment. Furthermore, intelligent PT based on DRL uses the function approximator as a neural network [29] to conduct PT. It is easy to encounter the high-dimensional discrete of the action space when applied to the actual scenario and complex network system, which makes DRL model tough to train and converge in the PT process. Taking the value-based Deep Q-Network (DQN) algorithm [10] as an example, the limitation of using DQN for automated PT lies in the complexity of the action space of different penetration scenarios. DQN evaluates all output actions' values, and selects the largest value as the best option in the entire action space. If the action space is large, we may encounter the problem that multiple actions have the same value [11]. PT is similar; the state and action space exponentially expand with the increasing hosts in different subnets. Consequently, applying the existing DRL algorithms to automate PT is difficult and unstable. Although in a relatively small network scene, it still faces the challenge of the state and action space explosion with hundreds or thousands of dimensions [42].

In summary, the existing PT methods are still challenged in three aspects: (1) Security experts play a significant role in the PT process, but over-reliance on security experts' manual operations and decisions will also increase the labor cost of PT; (2) When using the DRL algorithm for automating PT, it could encounter the problem of large state space and high-dimensional discrete action space, thus leading to convergence difficulty of PT training; (3) Most intelligent PT methods are verified in the virtual network environment instead of real penetration scenarios.

To address these challenges, for the first time, we introduce automatically-collected expert knowledge for agents to perform autonomous PT, proposing a general intelligent PT framework that incorporates Generative Adversarial Imitation Learning, denoted as GAIL-PT.We combine GAIL with RL / DRL to automate PT process, by modeling the penetration attacker as a agent. Specifically, due to imitation learning has achieved a great success for speeding up the convergence rate of reinforcement learning with learned strategy in prior, we adopt GAIL for accumulated expert knowledge for RL / DRL based PT. Then, the agent continuously interacts with the penetration environment through the guidance of expert knowledge, and the accumulated penetration experience will assist the agent in formulating new attack strategies with a higher penetration success rate adaptive to the dynamic environments. At last, for both small-scale and large-scale network scenes with high-dimensional discrete action space, GAIL-PT shows state-of-the-art (SOTA) penetration performance compared with baselines and also can be verified in practical scenarios.

The main contributions of this paper are as follows:

- Aiming at the problem that most PT methods over-reliance on security experts' manual operations and decisions, we construct an expert knowledge base and store state-action pairs as experiences

in different penetration scenarios so that the PT experts can take part in the decision-making in a lower-cost way.

- To solve the convergence difficulty problem of PT training and improve the penetration performance, we apply the GAIL network to PT based on RL or DRL for the first time. With the guidance of the expert knowledge, the agent is trained through the GAIL network to perform the action infinitely closed to the expert knowledge base, making the training process more stable and efficient for PT.

- We conduct extensive experiments not only in real target Metasploitable2 and small-scale networks with or without honeypot, but also for large-scale networks. The experiment results indicate that GAIL-PT shows the SOTA penetration performance in actual or simulated network scenes. In addition, more experiments also verified that the proposed GAIL-PT is a general leading framework suitable for different reinforcement learning methods, i.e., DRL, RL.

The organization of this paper is as follows. Section 2 provides a overview of penetration testing, penetration testing based on reinforcement learning and imitation learning. Section 3 supplies an overview of the reinforcement learning models we applied. Section 4 exhibits the outline of our method and demonstrates its details. Section 5 describes the experimental setup, then evaluates and discusses the results. Finally, in Section 6 the main findings and limitations of our work are summarized.

2. Related Work

In this section, we mainly present penetration testing, the current research status of penetration testing based on reinforcement learning and mainstream imitation learning methods.

2.1. Penetration testing

PT is a security exercise designed to evaluate the system's overall security by authorizing simulated cyberattacks on the computer system. Manual PT relies on the use of penetration tools. For instance, Nessus [49] and Nmap [25] use the target list to perform network vulnerability testing one by one without checking for real-time environment changes. They neither make any decisions automatically, nor work independently. As one of the best performing vulnerability exploits, Core Impact [15], establishes an attack plan automatically by security experts. However, its limitation lies in generating attack plans on the target environment before conducting PT. Consequently, Core Impact cannot be dynamically adaptive to the environment either. Another prevalent PT tool is Rapid 7 Nexpose [37]. It uses the vulnerability development, and penetration tool Metasploit [38] for vulnerability scanning and penetration. Although it can achieve flexibility by distributing multiple network scan engines, Nexpose is unable to learn attack policies. In short, most well-known commercial penetration testing tools show lots of advantages

but highly rely on the expertise and decision-making of PT experts. These tools suffer from the problem of poor adaptability, that is, once the environment changes, the attack strategy cannot be updated accordingly.

In addition to manual PT methods, automated PT methods have been further explored. Numerous automated PT methods define PT as a path planning problem for constructing attack graphs [35]. Cynthia and Swiler [32] firstly proposed a method based on the attack graph to analyze the system's vulnerability. The system divides the common attack database into atomic steps, distinct network configuration, topology details and attacker configuration files. The nodes and chain edges in the attack graph describe the different attack stages. The probability of success will be allocated to the chain edge, so different attack graph algorithms are used to identify the attack path with the highest probability of success. Based on the previous attack graph method, Kyle et al. [19] created the NetSPA attack graphics system; its innovation lies in allowing network defenders to assess security threats to select complementary strategies. NetSPA analyzes multiple targets in minutes using firewall rules and vulnerability scanning, drastically reducing attack graph building time. Considering the attack path's reliability, Xue Qiu et al. [34] proposed an algorithm for automatically generating attack graphs using Common Vulnerability Scoring System (CVSS) [27] for the first time to improve the trustworthiness of attack paths by eventually optimizing the network topology. As a result, once an attack graph is generated, the network topology is also optimized. However, most of these methods could simply output corresponding action sequences to handle static environments. In other words, they can go back and provide corresponding steps or guidelines for PT, but cannot perform PT dynamically and interactively in real PT scenarios. In summary, we concluded some details of the introduced manual and automatic PT tools as shown in Table 1.

Table 1: The details of manual and automatic PT tools

Category	Tool	Benefits	Drawbacks	Scenes	Integration
Manual PT	Nessus[49]	Vulnerability scanning precisely	Cannot learn attack strategies	Single server & Network topology	Low
	Nmap[25]				
	Nexpose[37]				
Automatic PT	Metasploit[38]	Semi-automated PT process	Highly rely on the expertise	Single server	High
	Core impact[15]	Generate attack plans automatically	Inability to update dynamically	Single server & Network topology	High
	Attack graph[32]	Attack sequence and effect are clear	Cannot interact dynamically	Network topology	Middle

2.2. Penetration testing based on reinforcement learning

In recent years, RL [54] has achieved remarkable success in game areas, e.g., Alpha Go [45], Open AI Five [4] and Alpha Star [52]. Particularly, in some games the agents based on RL have surpassed human players. Just like game rules, PT is also a dynamic decision-making process based on observing the environment. The agent in RL-based PT is trained to observe and explore the dynamic network environment, so as to learn the optimal policies through trial and error [6]. Intuitively, RL is an excellent option to stimulate and develop as an effective tool for penetration attacks against network security. Schwartz

and Kurniawati [42], Hu et al. [18], Zennaro and Erdodi [55] all describe PT as a POMDP, using RL to model the penetration environment and learn penetration strategies. However, the penetration agent merely observes the compromised host and the subnet, it does not observe the connection information of the global network. In other words, the agent learns the optimal strategy to execute the penetration attack through continuous trial and error learning. These studies show that RL can be applied to relatively small action spaces or CTF challenges.

Considering the observation state, the PT environment is different from the RL game environment, because the action space of the network penetration environment is discrete and high-dimensional. It is still a long-term challenge to extend DRL to a network with a more extensive action space [42]. Arnold et al. [7] proposed a framework called Wolpertinger, which uses the Actor-Critic (AC) framework to learn strategies in a large action space that combines sub-linear complexity. It could be unstable in a spare reward environment since the gradient cannot be back-propagated for the training of the actor network. To address the issue of sparse rewards in large action spaces, Zhou et al. [57] proposed an improved DQN algorithm NDSPI-DQN to optimize the penetration path. It effectively reduces the agents' action space by decoupling attack vectors, but the method is only suitable for simulation scenarios. Besides, Tran K et al. [51] proposed automated PT based on deep hierarchical reinforcement learning, denoted as HA-DRL, which uses algebraic action decomposition approach to deal with the sizeable discrete action space of PT simulators. However, the space of actions will grow up exponentially with the complexity of the network. Thus the training costs are relatively high. In addition, Bland et al. [5], Elderman et al. [9] and He et al. [14] proposed the application of multi-agent reinforcement learning for automated penetration testing of network security simulation scenarios, and they also proposed a learning algorithm of the optimal strategy. This method can achieve satisfying performance in simulated scenarios at the cost of high training complexity.

2.3. Imitation learning

Imitation learning [46] refers to automatic learning from examples provided by experts. The purpose of training the model is to fit the trajectory distribution of the strategy generated by the model with the trajectory distribution of the input. According to different ways of strategy optimization, imitation learning methods roughly divided into Behavioral Cloning (BC) [48][50], Inverse Reinforcement Learning (IRL) [3] and Generative Adversarial Imitation Learning (GAIL) [16].

BC was proposed by Torabi et la.[50] to directly learn the optimal action in the sampling state from the expert's empirical data without constructing a reward function and predict the corresponding optimal action in the new state after learning. Nevertheless, BC is prone to cascade errors. Based on BC, IRL finds a reward function that explains these strategies or behaviors by giving optimal strategies or trajectories. In order to improve the performance of IRL, Wang et al. [53] proposed Imitation Learning via Inverse Reinforcement Learning (IRL-IL) based on IRL. It is designed to solve the optimal strategy through RL, restore the expert strategy indirectly, and make a plan with long-term. Consequently,

it solves the cascade error problem of BC, leading to more substantial generalization and robustness. However, the linear reward function of most IRL-IL methods contains certain restrictions [24], i.e., the representation ability of its reward function is insufficient, and the set penalty item cannot be assigned to the expert strategy with a larger reward value as much as possible; what is more, the iterative solution of the RL sub-process requires lots of computing resources [17]. In order to make up for these limitations, Jonathan Ho et al. [16] combined the Generative Adversarial Network (GAN) [12] with IRL-IL proposed GAIL. The expert data is compared with the data generated by the agent network and optimized the model reversely so that the agent can learn the strategy of approaching the expert and alleviate the training problem of IRL-IL. In addition, the algorithm has achieved a considerable performance improvement beyond the current model-free methods when simulating problematic behaviors in large-scale high-dimensional environments [58].

3. Preliminaries

This section briefly introduces reinforcement learning and its algorithm Q-learning. Besides, we also introduce DRL algorithms based on different policies, i.e., Asynchronous Advantage Actor Critic (A3C) and Distributed Proximal Policy Optimization (DPPO).

3.1. Reinforcement learning

RL is a machine learning method based on the sequence interaction between the agent and the environment. RL is usually modeled as a Markov process to solve sequential decision-making problems. It can be represented by a four-tuple $< S, A, P, R >$, containing the state space S, action space A, state transition probability P and reward function R. Through continuous interaction with the environment, the agent a_t acts under the current state s_t according to the learned policy π. Simultaneously, the environment feeds back to the agent a scalar reward value $r(s_t, a_t)$ to evaluate the action's quality, and then transfers to the next state according to the state transition probability $P(s_{t+1} \mid s_t, a_t)$.

The objective of RL is to maximize the cumulative reward over time t, denoted by the reward function:

$$R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k} \tag{1}$$

where $\gamma \in [0, 1]$ is the discount factor used to measure the importance of current rewards to future rewards, the smaller the value, the agent only focus on the current rewards; on the contrary, the larger the value indicates that the agent will concentrate on to the future long-term returns.

In this paper, we consider the value-based RL algorithm Q-learning [29], the policy-based multi-threaded DRL algorithms A3C [28] and DPPO [41].

3.1.1. Q-learning

Q-learning is a representative algorithm in value-based model-free learning. Q represents the quality function $Q(s, a)$ of the policy π, which refers to the expectation when executing action a in the state of s at a specific moment. The algorithm associates all states s with action a to form a Q table to store the Q value of the state-action pair (s, a), and generates the best policy π^* by selecting the action with the highest Q value in each state. The update mechanism explaining in formula:

$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \tag{2}$$

where s, a are the current state and action, respectively. s' is the next state that emerges as a result of action a, a' is a possible action in state s'. r is the instant reward, α and γ represent learning rate and discount factor respectively.

3.1.2. Asynchronous advantage actor critic

Asynchronous Advantage Actor Critic (A3C) [28] is a DRL algorithm based on policy gradient. A3C adopts multi-thread mechanism, and uses Actor-Critic(AC) network structure in both of the chief network and the thread network. AC is divided into Actor network $\pi^{\theta'}(a|s)$ and Critic network $V^{\mu'}(s)$. The corresponding policy $\pi(a|s; \theta')$ is obtained by inputting the current state s, which indicates the probability of selecting an action a under the condition of state s and the Actor network parameters θ'.

In A3C, an advantage function $A(s, t)$ is constructed by using the output of the value function $V(s|\mu')$ to evaluate the policy adopted, μ' is Critic network parameters. When N-step sampling is used, the advantage function obtained is:

$$A(s, t) = r_t + \gamma r_{t+1} + \ldots + \gamma^{n-1} R_{t+n-1} \gamma^n V(s') - V(s) = R(t) - V(s) \tag{3}$$

where n means N-step, γ is the discount factor. $R(t)$ represents the reward function at the current time t, $R(.)$ represent the value of the current state s. Then calculate the first derivative of θ' in Actor network and the second derivative of μ' in Critic network to update the $d\theta$ and $d\mu$, respectively.

$$d\theta \leftarrow d\theta + \nabla_{\theta'} \log \pi(a|s; \theta') A(s|\mu') \tag{4}$$

$$d\mu \leftarrow d\mu + \partial A(s|\mu')^2 / \partial \mu' \tag{5}$$

3.1.3. Distributed proximal policy optimization

Similar to the A3C algorithm, the essence of Distributed Proximal Policy Optimization (DPPO) [41] algorithm uses the Proximal Policy Optimization (PPO) algorithm to train agent in a multi-thread distributed manner. The PPO algorithm is based on policy gradients and also uses a parameter θ of a neural network to approximate policy $\pi_\theta(a|s)$ directly. Besides, PPO is also implemented in an AC framework. PPO uses the advantage estimation function with reduced variance to stabilize the policy

gradient while learning the approximate policy $\pi_\theta(a|s)$. The most significant difference between the AC algorithm and PPO is that three networks are implementing the PPO algorithm: a Critic-network and two Actor-networks, just the old and new Actor-networks and the primary function is to limit the update step size of the Actor-network.

At each training step, PPO collects experience by executing the present policy within a set of time steps, calculates the practical returns and advantages, then utilizes batch learning to optimize a clipped surrogate objective L^{clip} that restrains the number of the updated strategies that may differ from the older strategies.

$$L^{clip}(\theta) = \mathbb{E}_t[min(r_t(\theta)\hat{A}_t, clip(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t] \tag{6}$$

where θ is the neural network parameter to approximate policy π. $\hat{A}_t$ refers to the estimation of the advantage function at the moment t. ϵ refers to the hyperparameter, used to restrict the change range of $r_t(\theta)$. It represents the probability ratio of the new policy to the old policy:

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{old}(a_t|s_t)} \tag{7}$$

In addition, the discrepancy between DPPO and A3C is that the update mechanisms of the chief and threads are different. The chief network of A3C does not participate in training. It only collects each thread's trained parameters, weighing them and sending them to each thread. The parameter of the child thread is new, and the parameter of the chief is old. On the contrary, the chief network of DPPO first starts training, sends its trained parameters to each thread, and then performs the process of parameter transfer and parameter update continuously, in which the threads always carry out the old strategy training and the main thread for the training of new policies.

4. Methodology

In this section, the method of automated penetration testing based on GAIL is introduced. First, the state-action pairs are automatically collected to construct an expert knowledge base when the pre-trained DRL / RL model executes successful post-exploit.Second, input the expert knowledge and the state-action pairs generated online by the different DRL / RL model into the discriminator of GAIL for training.At last, the discriminator's output reward is applied to guide the agent to perform the action with a higher penetration success rate to improve PT's performance. The overall framework of the proposed GAIL-PT is illustrated in Fig. 1, including three stages: ① Construction of penetration experts knowledge base; ② GAIL training; ③ Automate penetration testing based on GAIL.

4.1. Construction of penetration experts knowledge base

Expert sample data are necessarily acquired in advance when conducting PT based on GAIL to train agents in different PT environments. Different from general game scenes, e.g., Gym and Atari [10],

there are not any published expert samples for agent training in PT environment. Therefore, to our best knowledge, it is the first time, we construct penetration experts knowledge base for PT.

There are several ways to obtain PT expert knowledge samples, such as saving penetration rules defined by network security experts. It is usually difficult and expensive to collect these rules manually. Therefore, when using DRL or RL model to perform PT, we automatically collect the state-action pairs of the agent when successfully executing post-exploit, then save them into the expert knowledge base constructed online. At present, DeepExploit [20] is a better way to exploit the single target host Metasploitable2 [36] than manually using Metasploit.Using DeepExploit to perform automated PT will encounter three different results in an authentic single target host: unsuccessful exploit, successful exploit and successful post-exploit.
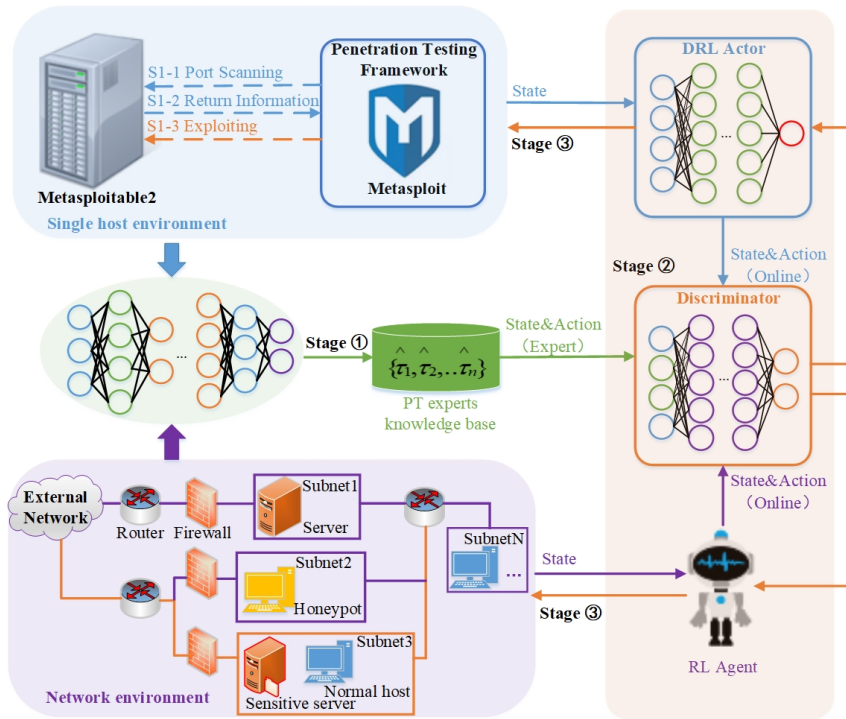

Figure 1: Overall framework of GAIL-PT. Stage ① Construction of penetration experts knowledge base: collecting PT expert samples in different penetration scenarios, storing the state-action pairs when using RL / DRL model executing successful exploitation to construct the expert knowledge database. Stage ② GAIL training: Input both the expert samples and state-action pairs generated by the different RL / DRL model online into the discriminator in GAIL for training. Stage③ Automate penetration testing based on GAIL: RL / DRL agent applies the new action from the well-trained discriminator to automate the PT process and execute more effective penetration attacks.

As shown in Fig. 2, when the Kali server initiates a penetration attack to the target host, if stage 1 executes successfully, the target server will establish a session connection with the Kali server at the same time like stage 2, then will use the payload to execute related operations on the target host to gain the root privilege. We call the process of privilege promotion post-exploit, which corresponds to stage 3. Suppose stage1, stage 2 and stage 3 all execute successfully, in that case, it indicates that the Kali server

has successfully exploited the target host and performed subsequent privilege promotion operations by the payload, which means the post-exploit is also successful. We set different reward values for mixed penetration results and set the highest value for a successful post-exploit. Therefore, in the process of PT, the reward value obtained by the agent can be used to confirm the penetration result at that moment. Besides, We collected the state-action pairs with the highest reward value as the expert sample data and stored them into expert knowledge base in the form of one-to-one correspondence between states and actions.
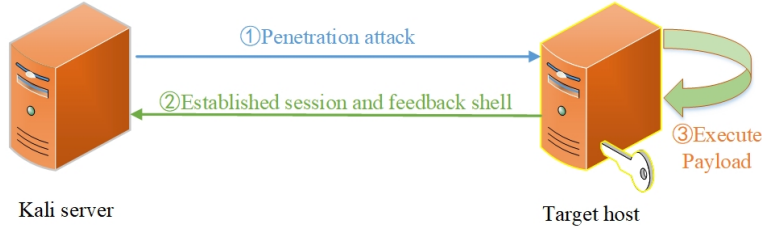

Figure 2: Process of executing integral penetration attack. Stage ①: Kali server launches penetration attack toward Target host. Stage ②: Target host establishes session and feedback shell with Kail server after successful exploit. Stage ③: Kali server uses the payload remotely to the target Server for executing related operations to promote the privileges and obtain its root authority.

Like the single target host scenario, the mode of collecting expert sample data in the network scenario also means storing the state-action pairs with the highest reward values when exploiting the target sensitive hosts successfully at the lowest cost. Generally speaking, construction of penetration experts knowledge base is similar for both exploiting a single target host and exploiting a network. Specifically, exploiting a host means acting some related operations on the target host, like executing the payload, etc., finally gaining control of the target host. In the network scenario, after gaining the control of one host, the attacker can use the compromised host as a stepping stone according to network connection relationships to carry out lateral movements for obtaining control of the final sensitive host. In other words, exploiting a network is just like finding an optimal attack path, and the path shows the exploit relationship between some hosts in different subnets. It is worth mentioning that the most significant distinction between our approach and the attack graph lies in our method does not need to know the network topology information in advance but finds an optimal attack path through continuous trial and error training.

For better understanding, an illustration example is shown in Fig. 3. There are various combinations of state-action pairs due to the diversity of vulnerabilities when exploiting the real target host Metasploitable2. The five dimension state represents operating system, port service, product version, protocol and exploit module type. The action corresponds to the number of payloads in Metasploit. Take the first row of data of a single target host as an example; 0.875 means the operating system is Linux, the port service is SSH (Secure Shell), the product version is 0.0, the protocol is TCP (Transmission Control Protocol), and the exploit module type is 0 represents choosing the first item of the exploit list. As we

11

know, there are 593 (0 ∼ 592) payloads in Metasploit. Action 53 means choosing the 53rd payload to execute the exploit operation.

However, for expert sample data in the network scenario, its state represents the configuration details and vulnerability information of all hosts observed by the agent in the different network environments. Besides, action means the operations of vulnerability scanning, exploit and privilege promotion executed by the agent arriving at the target host. Take the small-scale scenario's state and action as an example. The state vector consists of different host vectors for each subnet and starts with the first host. At the beginning of the state vector, the host location is represented by a 10 bit one-hot encoding. The numbers, i.e., 0 means false and l means true, behind successively represent compromised, reachable, discovered, value (specific value of different hosts), discovery value, access, os (operating system, i.e., 1 means Windows, otherwise means Linux), exploit services and privilege promotion processes. Besides, the additional final row in the state vector, representing action success, connection error, permission error and undefined error. It also uses 1 or 0 to represent true or false.

There are eight hosts in the small-scale network scenario (the scenario is shown in Fig.9), including two types of the operating system: Windows and Linux, three types of exploit services: SSH, HTTP (Hyper Text Transfer Protocol), and FTP (File Transfer Protocol), two types of privilege promotion process: Tomcat and daclsvc, two types of firewall restrictiveness: HTTP and SSH. Consequently, the actions could have 72 (8×9) types of operation combinations. The 16th action means executing privilege promotion by the Tomcat process on the sensitive hosts in subnet 2, which is critical in exploiting the small network scenario. In short, the state-action pair corresponding to the best penetration attack path, and is unique and deterministic, so the expert sample data is relatively fixed. After that, we only did simple and repeated expansions of the sample size.

Single target host	State	Action
	[0.875, -0.29729729729729, 0.0, -0.15737473535638674, 0]	53
	[0.875, -0.5135135135135, 1.31, -0.07551164431898377, 0]	62
Metasploitable2	[0.875, -0.40540540540543, 8.37, -0.6146788990825688, 0]	224
	[0.875, -0.35135135135137, 3.3, -0.13196894848270996, 0]	466
	...	...
	[0.875, -0.4594594594595, 5.051, -0.1460832745236415, 0]	472

Network scenario	State	Action
Small-scale	[0. 1. 0. 0. 0. 1. 0. 0. 0. 0. 1. 1. 1. 0. 0. 1. 1. 0. 0. 0. 1. 0. 0. 0. 0. 1. 0. 0. 1. 0. 0. 0. 0. 1. 1. 1. 100. 0. ···. 0. 0. 100. 0. 0. 1. 0. 0. ···.0.]	16
Small-scale with honeypot	[0. 1. 0. 0. 0. 1. 0. 0. 0. 0. 1. 1. 1. 0. 0. 1. 1. 0. 0. 0. 1. 0. 0. 0. 0. 1. 0. 0. 0. 0. 1. 1. 1. 100. 0. ···. 1. 1. -100. 0. 0. 0. 1. 0. ···.0.]	16
Large-scale	[0. 1. 0. 0. 0. 0. 0. 0. 1. 0. 0. 0. 0. 1. 1. 1. 1. 1. 2. 0. 0. 1. 1. 0. 0. 0. 1. 0. 0. 0. 0. ···. 0. 0. 0. 0. 1. 1. 1. 100. 1. 2. 1. 0.···. 1. 2. 1. 0. 0. 1. 0. 0. 1. 0. 1. 1. ···. 0. 0. 0.0. 0. 1. 1. 0. 0. 0. 1. 1. 1. 100. 1. 2. 0. 1. 0. 0. ···. 0. 0.]	315

Figure 3: Examples of PT expert knowledge base.

4.2. GAIL training

After collecting the PT expert sample data, this section introduces training the GAIL network.

Our proposed GAIL contains three different neural networks: Policy Neural Network (Actor), Value Neural Network (Critic) and Discriminator Network (Discriminator). The Actor consists of a four-layer network. The hidden layers comprise three dense layers with output dimensions of 50, 100 and 200 connected through the ReLU activation function. The last is the output layer with a Softmax activation function, and the output dimension corresponds to the number of actions.

The network structure of Critic is similar to that of Actor. The only difference is that the last output layer is linear, the dimension is one. Discriminator, known as Reward Neural Network, has a four-layer network with three hidden layers of the same structure as the Actor's hidden layers. The last output layer is activated by Sigmoid and outputs a specific probability value as a discounted reward.

The process of GAIL-training is exhibited in Fig. 4. The Actor, Critic and Discriminator are alternately trained. The inputs of the Actor and Critic are both currently observed states, while the output of Actor is the corresponding action, and the output of Critic is the value of the current state. The input of the Discriminator are the state-action pairs of the agent and PT experts. Besides, the state-action pairs of the expert sample are only used for training, and the output of the discriminator is the discount reward value.
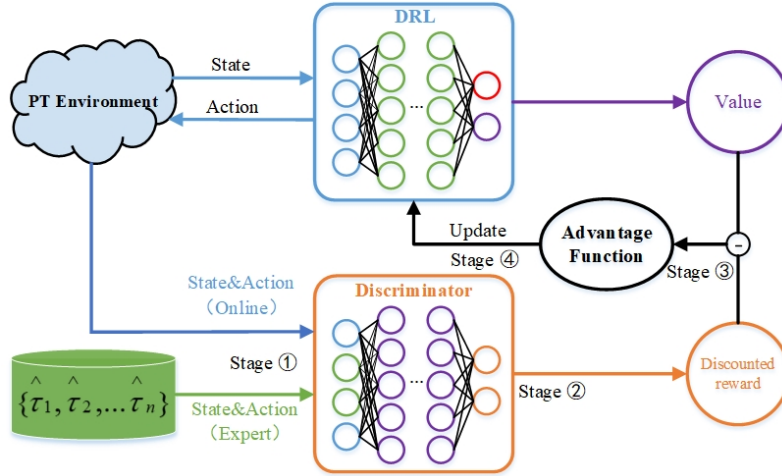
Figure 4: The Process of GAIL training. Stage ①: Put the state-action pairs from PT expert knowledge base and generating by online model respectively into the Discriminator for training together. Stage ②: Through maximizing the standard action reward given by the sample of PT experts and minimizing the action reward output by the agent to train the Discriminator. Stage ③: Subtract the discount reward and the value output by the Discriminator and the different DRL model to get the advantage function. Stage ④: Use advantage function to update DRL model meanwhile guiding agent output actions with a higher probability of successful penetration.

The main processes of GAIL-training are as follows: (1) Put the state-action pairs from PT expert knowledge base and generating by DRL / RL online model respectively into the Discriminator for training together. (2) Through maximizing the standard action reward given by the sample of PT experts and

minimizing the action reward output by the agent to train discriminator. (3) Use the discriminator's output to replace the original model's reward function to guide the training directions. (4) The DRL model output predicted distribution is infinitely close to the trajectory of PT expert knowledge base as much as possible.

During the training process of imitating PT experts strategy, the Actor network is used to replace the generator G, its output action and state are paired into the Discriminator, and compared with the expert data. The output of the Discriminator $D : S \times A \rightarrow (0,1)$ is considered as the reward signal to guide strategy learning in imitation learning. Therefore, regarding the training of policy and reward function as the process of a game that is analogous to the game of G and D in Generative Adversarial Network (GAN), the objective function $L_{GAIL}(\pi, D)$ as follows:

$$\min_{\pi} \max_{D} L_{GAIL}(\mathbb{E}_{\pi}[\log D(s,a)] + \mathbb{E}_{\pi_E}[\log(1 - D(s,a))] \tag{8}$$

The π in the formula represents the policy that should be learned, and π_E means the expert strategy. The first term $\log D(s,a)$ denotes the judgment of the Discriminator on the factual data, and the second term $\log(1 - D(s,a))$ denotes the judgment of the generated data. Through such a maximum-minimum game process, G and D are optimized alternately to train the Actor and Discriminator networks. In the training process, we get the total loss by adding the expert loss and the agent loss, minimizing the loss function by gradient derivation to update the Discriminator and Actor network parameters in reverse:

$$Loss = \min_{D \in (0,1)^{S \times A}} \sup \mathbb{E}_{\pi}[\nabla_{\omega} \log D_{\omega}(s,a)] + \mathbb{E}_{\pi_E}[\nabla_{\omega} \log(1 - D_{\omega}(s,a))] - \lambda \nabla_{\theta} H(\pi_{\theta}) \tag{9}$$

where $H(\pi_{\theta}) \triangleq \mathbb{E}_{\pi}[-\log \pi_{\theta}(a|s)]$, it represents the entropy of the imitation policy, controlled by a constant $\lambda(\lambda \geq 0)$, as the policy regular term in the loss function, and ω and θ signifies the parameter of Discriminator and Actor network separately.

4.3. Automated penetration testing based on GAIL

After GAIL's training is completed, automated penetration testing can be performed based on GAIL. This section narrates the automated penetration testing process based on GAIL. The training framework has been shown in Fig. 1.

4.3.1. Automated penetration testing based on GAIL combined with DRL

In the actual scenario of exploiting the single target Metasploitable2, we define the PT process as an MDP at first. The states s of the agent is defined by the information obtained by Nmap through port scanning; secondly, we define the actions a selected by the agent in the PT, the actions corresponding to the payload sequences list in Metasploit Framework (MSF); at last, the reward r is set according to the penetration result of the applied payload. We use four different DRL models to train the agent, such as A3C, DPPO, A3C-GAIL and DPPO-GAIL. What is more, we all used a multi-threaded mechanism for training to speed up the learning efficiency of the agent.

14

Each of the A3C and DPPO models we used consisted of 4 layers neural networks: the input layer, the hidden layers with three fully connected and the output layers. The state inputs of the model, i.e., operating system, port service, product version, protocol and exploit module type, are all obtained by scanning the target through the vulnerability scanning tool Nmap. The action output of the model corresponds to the probability distribution and the state value of each payload in Metasploit. For the Actor network in the model, we use Softmax as the activation function, and the output of the Actor network is the probability distribution $P(a)$ of each action; while the Critic network uses Linear as the activation function, and its output is the state value $V(s)$ corresponding to the current state.

We introduced a discriminator to participate in the Actor and the Critic network training during the Actor and the Critic network training. There are specific differences from A3C and DPPO models in the training process based on the DRL model combined with GAIL. The difference between the value function and the discriminator output discount reward is applied as the advantage function to guide the training of the policy network, and the output of the GAIL model is consistent with the original model.

The DRL model is regarded as an attacker in the learning process. According to the input vulnerability information, the agent will choose an effective payload to exploit the target host until further getting its root authority. Therefore, obtaining the root authority of the target host is a sign of a successful exploit. To be more precise, it is the post-exploit after the privilege promotion operation has been done, and then the agent will be rewarded. In addition, the value of rewards r depends on whether the PT attack is effective.

As shown in Fig. 2, if the agent executes a penetration attack from the Kali server to the target host, establishing a session connection with the Kali server and then executing the payload to gain root authority of the target host favorably. It indicates that the post-exploit is prospering, and the agent will obtain the highest reward. On the contrary, if the agent only initiates the penetration attack without establishing a connection with the Kali server, the agent will be given a lower reward at this time. Furthermore, if the agent has not exploited the target host successfully, its reward will be negative. The specific settings of the reward are as follows:

- The post-exploit is successful, r=100;

- The exploit is successful, but the root authority for target host is not obtained, r=1;

- The exploit is failure, r=−1;

- For other operations, regardless of whether the exploitation is successful or not, r=−1.

Negative reward r=−1 means to punish the agent, and the agent tries to maximize the reward all the time. Therefore, if the agent has been punished with negative rewards, it will try to reach the penetration goal as soon as possible. It corresponds to a situation where an actual attacker realizes detecting vulnerabilities and exploiting them successfully. In addition, as long as the agent gets a reward, it symbolizes that each round of training is over.

Details regarding the automated PT via GAIL combined with DPPO are presented in **Algorithm 1**.

Algorithm 1 DPPO-GAIL

Input: Expert PT trajectories $\hat{\tau}_{GAIL} \sim \pi_E$, five-dimension state s including operating system type, opening-port service name, product version, penetration module type, and key information of exploiting target type from using Namp to scan Metaspolitable2, iterations $i = 200000$.

Output: Action a means payload which is infinitely close to the expert samples a_E.

1: Initialize DPPO actor-net and GAIL discriminator-net parameters θ_0; i_0.

2: **for** $i = 0, 1, 2, \ldots$ **do**

3:      Generate state-action pairs from DPPO model online and store them as π_t.

4:      Sample online trajectories $\hat{\tau}_i \sim \pi_{\theta_i}$.

5:      Update the discriminator-net parameters from ω_i to ω_{i+1} with the gradient:

6:      $G_d = \hat{\mathbb{E}}_{\pi_{GAIL}}[\nabla_\omega \log D_\omega(s, a)] + \hat{\mathbb{E}}_{\pi_t}[\nabla_\omega \log(1 - D_\omega(s, a))]$.

7:      From DPPO actor-net take a policy step from θ_i to θ_{i+1} respectively.

8:      Using the TRPO rule with the cost function $\log(D_{\omega_{i+1}}(s, a))$.

9:      Specifically, DPPO takes a L^{clip} natural gradient step to update old and new policy with (6), and:

10:      $\hat{\mathbb{E}}_{\pi_{GAIL}}[\nabla_\theta \log \pi_\theta(a|s)Q(s, a) - \lambda \nabla_\theta H(\pi_\theta)]$,

11:      where $Q(\bar{s}, \bar{a}) = \hat{\mathbb{E}}_{\pi_{GAIL}} \left[\log \left(D_{\omega_{i+1}}(s, a) \right) \mid s_0 = \bar{s}, a_0 = \bar{a} \right]$

12: **end for**

4.3.2. Automated penetration testing based on GAIL combined with RL

When conducting penetration testing in the network attack simulator NASim[44], we also model the PT process as an MDP at first, the state is defined as the configuration details and vulnerability information of all hosts observed by the agent in the different network environments. The larger the number of hosts, the larger the state scale. Besides, action means the manipulations of vulnerability scanning, exploit and privilege promotion executed by the agent arriving at the target host. We defined the action as an attack vector, $< m, c >$ means the agent manipulates m on the host computer c. In each training step, the scale of operating space for an agent reaches $(P \times Q)$, where P are the number of hosts in the network, Q represents the number of manipulations that exploit the final target host could perform. After the agent executes a particular operation, the environment will feedback a corresponding reward value. The reward value is calculated as shown in (10), including the value of the host and the cost of exploits, privilege promotion, or scanning operations on the host.

$$R = \sum_{c \in C} value(c) - \sum_{m \in M} cost(m) \tag{10}$$

We use the RL algorithm Q-learning as a benchmark to automatically generate penetration paths and optimize its process for path-finding. During the training process, the agent will select the corresponding representative vulnerabilities according to the service to calculate the vulnerabilities' probability. C in

Equation (11) represents the set of host computers in different networks that the agent can exploit, and M represents the set of manipulations the agent can adopt. The attacker intends to maximize the cumulative reward and explore the optimal penetration path, until exploiting the most valuable sensitive target hosts successfully with as few operations as possible.

$$\max_{\pi} \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_t + 1)|\pi] \tag{11}$$

Subsequently, we combined GAIL with the Q-learning algorithm. The Q-learning-GAIL algorithm is applied to three different scenarios of a small-scale network (considering whether there exists a honeypot) and a large-scale network to train the agent to accomplish the optimization of finding the penetration path. The Q-learning algorithm is different from the DRL algorithm. Its optimization process is to continuously minimize the difference $TD - error$ between the target value Q_{tar} and the realistic evaluation value Q_{eva}. When calculating the current value, Q_{eva}, the immediate reward r needs to be are summed with the discounted maximum target value Q_{tarnex} in the next state. Therefore, to improve the optimization effect, we superimpose the instant reward r and the reward r_d of the discriminator output, which is used to guide the learning of the imitation policy, so that leads the action distribution of the RL model is close to PT experts samples as much as possible.

Details regarding the automated PT with GAIL combined with Q-learning are presented in **Algorithm 2**.

Algorithm 2 Q-learning-GAIL

Input: Expert PT trajectories $\hat{\tau}_{GAIL} \sim \pi_E$, states s including configuration knowledge and vulnerability information of all hosts, iterations $i = 20000$.

Output: Actions a means exploits or privilege escalation which are infinitely close to the expert samples a_E.

1: Initialize $Q(s, a)$ arbitrarily and GAIL discriminator-net parameters ω_0.
2: **for** $i = 0, 1, 2, \ldots$ **do**
3:      Generate state-action pairs from RL model online and store them as π_t.
4:      Sample online trajectories $\hat{\tau}_i \sim \pi_{\theta_i}$, initial s..
5:      Update the discriminator-net parameters from ω_i to ω_{i+1} with the gradient:
6:      $G_d = \hat{\mathbb{E}}_{\pi_{GAIL}}[\nabla_\omega \log D_\omega(s, a)] + \hat{\mathbb{E}}_{\pi_t}[\nabla_\omega \log(1 - D_\omega(s, a))]$.
7:      Choose a from s using derived from $Q(\epsilon\text{-greedy})$.
8:      Using the TRPO rule with the cost function $\log(D_{\omega_i+1}(s, a))$.
9:      Take action a, observe $r, \overline{s}$.
10:      From discriminator-net get discounted reward r_d.
11:      Update $Q(s, a) \leftarrow Q(s, a) + \alpha[r + r_d + \gamma \max_{\overline{a}} Q(\overline{s}, \overline{a}) - Q(s, a)], s \leftarrow \overline{s}$.
12: **end for**

4.4. Convergence analysis of GAIL

Before GAIL [16] was proposed, Ho et al. [24] first proposed the inverse reinforcement imitation learning (IRL-IL) framework:

$$\max_{\pi} \min_{r} \mathbb{E}_{\pi}[r(s,a)] - \mathbb{E}_{\pi_E}[r(s,a)] + \psi(r) \tag{12}$$

where $r(s,a)$ represents the reward function, and the reward r corresponding to the state-action pair (s,a), π_E represents expert strategy, π represents the strategy need to be learned, $\mathbb{E}_{\pi}[\cdot]$ denotes the expectations about the strategy, and $\psi(r)$ is the penalty term of the reward function. Based on this framework, GAIL gave a particular but more reasonable form of penalty term:

$$\psi_{GAIL}(r) \triangleq \begin{cases} \mathbb{E}_{\pi_E}[g(r(s,a))] & \text{if } r > 0 \\ +\infty & \text{otherwise} \end{cases} \quad g(x) = \begin{cases} x + \log(1 - e^{-x}) & \text{if } x > 0 \\ +\infty & \text{otherwise} \end{cases} \tag{13}$$

The valuable feature for the penalty term $\psi_{GAIL}(r)$ is that it encourages the new reward function $g(x)$ to assign a more considerable reward value to the expert strategy π_E ; its special feature is that if the reward function satisfies a specific form:

$$r(s,a) = -\log D(s,a) \tag{14}$$

where $D(s,a)$ represents the probability that the Discriminator D discriminates (s,a) produced by the expert strategy. It happens to combine IRL-IL with Generative Adversarial Networks (GANs): the Actor network of inputting state and generating policy can be compared to Generator G. The reward function of inputting state-action pairs and outputting discount rewards can be compared to Discriminator D. The optimization process of the policy based on the current reward function can be analogous to the training process of the Generator. Similarly, the optimization process of the reward function can be analogous to the training process of the Discriminator.

Therefore, the core of GAIL is to minimize the Jensen-Shannon divergence between the state-action sample distribution $\rho_{\pi}(s,a)$ generated by the strategy π and the expert sample distribution $\rho_{\pi_E}(s,a)$:

$$\psi_{GAIL}^{*} = \sup_{D \in (0,1)^{S \times A}} \mathbb{E}_{\pi}[\log D(s,a)] + \mathbb{E}_{\pi_E}[log(1 - D(s,a))]. \tag{15}$$

Equation (15) proportional to the best negative logarithm loss of the binary classification problem of distinguishing the different state-action pairs. The results show that the optimal loss is the Jensen-Shannon divergence under a constant offset and scale, which is also the square measure between the standardized occupancy distributions:

$$D_{JS}(\overline{\rho}_{\pi}, \overline{\rho}_{\pi_E}) \triangleq D_{KL}(\overline{\rho}_{\pi}||(\overline{\rho}_{\pi} + \overline{\rho}_{E})/2) + D_{KL}(\overline{\rho}_{E}||(\overline{\rho}_{\pi} + \overline{\rho}_{E})/2) \tag{16}$$

where $\overline{\rho}_{\pi} = (1-\gamma)\rho_{\pi}, \overline{\rho}_{\pi_E} = (1-\gamma)\rho_{\pi_E}$. Regarding the causal entropy as a policy regular item, it is controlled by $\gamma(\gamma \geq 0)$, and in order to make the classification clearer, abandoning the standardization of the occupancy measurement $1 - \gamma$, thereby obtaining a new imitation learning algorithm:

$$\min_{\pi} \psi_{GAIL}^{*}(\rho_{\pi} - \rho_{\pi_E}) - \lambda H(\pi) = D_{JS}(\rho_{\pi}, \rho_{\pi_E}) - \lambda H(\pi) \tag{17}$$

So we can get a strategy, and its occupancy rate indicator minimizes the Jensen-Shannon divergence between the generated and expert samples. The above Equation (17) minimizes the accurate metric between the occupancy metrics. Therefore, unlike the linear apprentice learning algorithm, it can imitate expert strategies accurately. It also proves that the state-action pair generated by the instant strategy π can fit the sample distribution of the expert strategy π_E finally by minimizing the difference between $\rho_\pi(s,a)$ and $\rho_{\pi_E}(s,a)$ so that GAIL could come to converge.

4.5. Analysis of algorithm complexity

In the realistic scenario of exploiting Metasploitable2, the automated PT process is completed based on A3C-GAIL and DPPO-GAIL. As for DRL algorithms, the algorithm's complexity is related to the number of training episodes m and maximum steps t per episode. For the A3C algorithm, the thread network and the chief network adopt AC structure with only one for-loop statement nested. Besides, the GAIL network also nests one for-loop. Because of the GAIL, the Actor and the Critic adopt an alternative train mechanism. After combining with the GAIL, the algorithm complexity of A3C-GAIL is $O(n_m n_t)$. In addition, the DPPO algorithm uses the new and old policy update mechanisms. The training program of the chief network contains three for-loop statements, two of which are nested in another for-loop in a juxtaposed relationship. One of these two for-loop statements are used for model training. The other is used for agent choosing action. Then combined with GAIL, the algorithm time complexity of DPPO-GAIL is $O(n_m n_t)$.

The Q-learning algorithm's complexity is relatively small for the network scene to automatically find and optimize penetration paths because there is only one for-loop statement in it. Consequently, after nesting with GAIL, the algorithm time complexity of Q-learning-GAIL stays at $O(n)$.

5. Experiments

The experiments we conducted are divided into the following two parts. First, we test the penetration performance of using the combination of GAIL network and DRL algorithms A3C-GAIL, and DPPO-GAIL to automate PT in the real Metasploitable2 scenario. Then, considering the complexity of the actual penetration environment, we employ the network attack simulator NASim to find the optimal penetration path automatically. To testify the effectiveness of the RL model, we combine GAIL network with RL algorithm Q-learning, and evaluate the penetration performance of using the Q-learning-GAIL in a small-scale network (considering whether there exists honeypot [47] [8]) and a large-scale network.

Specifically, all experiments will perform on a single server with an Intel i7-7700K CPU running at 4.20GHz, 64 GB DDR4 memory, 4 TB HDD and 2 TITAN Xp 12 GB GPU card, operation system is Ubuntu 16.04, and employ Python 3.6 with tensorflow-2.1.0 and Python 3.7 with torch-1.5.0 for exploiting in metasploitable2 and NASim.

5.1. Research questions

The experiment will be analyzed based on different experimental scenarios around the following four research questions (RQs):

RQ1: Can expanding expert knowledge samples improve penetration performance in exploiting Metasploitable2? What is the effect of the different amounts?

RQ2: How efficient is GAIL's penetration when exploiting Metasploitable2? Can GAIL achieve the SOTA performance?

RQ3: Does expanding the expert sample amount still improve the penetration performance in different network scenarios?

RQ4: Can GAIL also achieve the SOTA penetration performance in different network scenarios?

5.2. Single host senario

5.2.1. Explanation for Metasploitable2

We used Metasploitable2 as the authentic single target host. The Metasploitable2 equipped an elaborated Ubuntu operating system designed as a security tool to test and demonstrate the environment for common vulnerabilities and attacks. Its primary purpose is to use as a target for launching penetration attacks by MSF. The target host opened 23 ports, and some of them opened port vulnerability information is shown in Table 2, including some high-risk ports such as 21, 23, 8180, etc. Many unpatched high-risk vulnerabilities include Samba, MSRPC, Shell and Command injection vulnerabilities, etc. Besides, some vulnerabilities open various services to the outside environment, and the database allows external connections. The user passwords in the system are vulnerable, and it equips with web vulnerability drill platforms such as DVWA and Mutillidae.

Table 2: Part of the opened port information of Metasploitable2

Port	State	Service
21/tcp	open	FTP
22/tcp	open	SSH
23/tcp	open	Telnet
…	…	…
8180/tcp	open	HTTP

5.2.2. Experiment setting

In order to evaluate the penetration performance of different algorithms in exploiting the target Metasploitable2, we first choose DeepExploit as the benchmark, which automates PT based on the A3C algorithm.Second, DPPO is different from A3C only in the update mechanism of the thread and chief network, so we also explored the feasibility of using the distributed algorithm DPPO to automate the

PT process, and tested its performance on exploiting the same target.At last, we combined the GAIL network with A3C and DPPO algorithm, respectively. We Applied the A3C-GAIL and DPPO-GAIL algorithms to automate the PT process, exploited the target host Metasploitbale2, and tested its penetration performance.

This section uses the above four different algorithm models to automate PT, and the ultimate goal is to execute the post-exploit successfully. The four algorithms all use 20 threads for 200,000 episodes of training and the same hyper-parameter value settings, details are shown in Table 3.

Table 3: Hyper-parameter values of four algorithms.

Hyper-parameter	Value
Max steps per episode	30
Max update step size per episode	5
Learning rate	0.0001
Greedy rate	0.8
Entropy coefficient	0.01
Loss coefficient of Critic-net	0.05

5.2.3. Evaluation metrics

In the single target host scenario, we apply the following five indicators from the perspectives of DRL performance and penetration performance to analyze the answers of ***RQ1*** and ***RQ2***:

- **Total reward.** For the training of DRL models, the reward value most intuitively reflects the learning quality of the agent, so we apply the total reward as the first indicator. The total reward is summed in every episode, which can be used to evaluate the overall performance of the algorithm;

- **Loss value.** For the training of DRL models, we can judge the stability of the model training by the loss value, so we apply the loss value as the second indicator. The loss value is calculated when the chief and child network parameters are updated, which can also be used to evaluate the overall performance of the algorithm;

- **Post-exploit count.** In the process of exploiting the Metasploitable2, post-exploit means the agent exploited the vulnerabilities service in a different open port and successfully executed the payload to gain the privilege for the target host. The more successful post-exploit, the more effective the PT is. It is the most intuitive reflection of penetration performance;

- **Time cost.** In the process of exploiting the Metasploitable2, time cost means the time spend in 200000 episodes training for using DRL models to exploit the target host. It reflects the cost we should pay. The less the time cost, the better the penetration performance is;

- **Time cost under exploiting limited vulnerabilities.** Considering the limitations of applying the MSF to exploit the target host Metasploitable2, no matter using the MSF manually or automatically to exploit Metasploitable2, we could only exploit the vulnerabilities of 11 types of services: VNC (Virtual Network Computing), Telnet, SSH, RPC (Remote Procedure Call), ProFTPd, PostgreSQL (Postgre Structured Query Language), PHP (Hypertext Preprocessor), MySQL (My Postgre Structured Query Language), IRC (Internet Relay Chat), HTTP (Hypertext Transfer Protocol), Apache. Therefore, we compared the time to exploit these vulnerabilities and used it as the last indicator to evaluate the penetration performance.

5.2.4. Research question 1

Can expanding expert knowledge samples improve penetration performance in exploiting Metasploitable2? What is the effect of the different amounts?

We will answer the question through analyzing the experiment results from the two indicators: Post-exploit count and Time cost.

This section used DeepExploit as the benchmark and then built an expert knowledge base by collecting expert state-action pairs when the post-exploit is successful. We tested the count of post-exploit and time cost of applying A3C-GAIL performing automate PT with 10,000, 25,000, 40,000, 55,000 and 70,000 pairs of five different expert sample amounts in the same training episodes, using the two evaluation indicators to overview the impact of the expanding expert samples on the penetration performance.
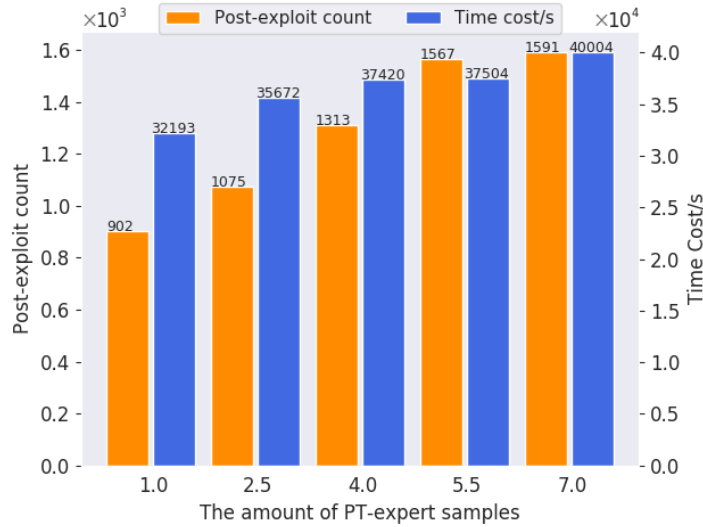
Figure 5: The impact of different expert sample sizes on penetration performance in exploiting the target host Metasploitable2

It can be observed from the experiment results in Fig. 5 that under the premise of considering the count of successful post-exploit and time cost comprehensively, the increase in the sample size of experts from 10,000 to 55,000 pairs, the improvement in penetration performance positively correlating with the
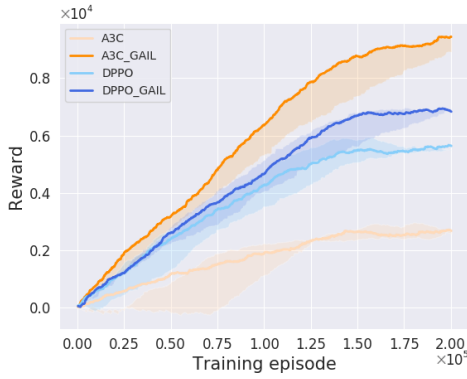
(a) Comparison of reward
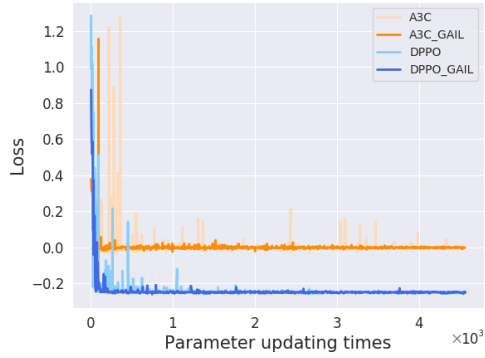
(b) Comparison of loss value

Figure 6: Different algorithms compared the reward value (left) and loss value (right). The training episodes of the four algorithms are all 200,000, they are all trained five times, and we present the best results of three of them in the figure. The shaded part of the left figure represents the fluctuation of the reward value in each episode, and the solid line represents the most stable training result. The loss value on the right figure corresponds to the situation when the reward value is the most stable.

size of the expert samples. When the expert sample size is expanded to 70,000 pairs, although the count of successful post-exploit increases slightly, the time cost is higher. Therefore, we determined that the sample size is around 55,000 pairs to achieve better penetration performance. Besides, we also use the DPPO-GAIL model with this expert sample size to automate PT in the subsequent experiment, and the final expert sample size is 57296 pairs.

> ***Answer to RQ1***: Penetration performance is improved by expanding the amount of expert samples in the range of 55,000 pairs, but it tends to be saturated when the expert samples exceed this range.

5.2.5. Research question 2

> How efficient is GAIL's penetration when exploiting Metasploitable2? Can GAIL achieve the SOTA performance?

We will answer the question through analyzing the experiment results from the above proposed five indicators: Total rewards, Loss value, Post-exploit count, Time cost, Time cost under exploiting limited vulnerabilities.

Total rewards. As shown in Fig. 6(a), at the beginning of training, the total reward value obtained by the agent is relatively low, and the loss function has not yet converged. With the continuous learning of the agent, when the training reaches 150,000 rounds, the reward function values for four algorithms all show a trend of convergence. Although the introduction of the GAIL mechanism does not speed up the convergence of the reward function of the A3C and DPPO models, the total reward value of the A3C-GAIL and DPP-GAIL are higher than the original DRL model, especially the total reward value of the A3C-GAIL is 243% higher than DeepExploit.

Loss value. In the course of many experiments, we found that the loss value of the benchmark DeepExploit will explode in the later stage of training, which causes the model to fail to convergence.

Therefore, we modify the loss function in the original DeepExploit algorithm and use modified Deep-Exploit as a new benchmark to compare with the other three algorithms. Since the multi-thread training mechanism only calculates the loss value of the model when passing the parameters between the thread and the chief network, the meaning of the ordinate is not equal to the number of training episodes. From the convergence of the loss function of the four algorithm models in Fig. 6(b), we can see that after updating the DRL algorithm model with GAIL mechanism 10,000 times, the fluctuation of the loss value should be stable to the original DRL model, particularly the loss fluctuation contrast between A3C-GAIL and A3C is more evident than DPPO-GAIL and DPPO.

Post-exploit count & Time cost. In addition to the analysis of the total reward and loss value from the perspective of the deep reinforcement learning model, we also compare the penetration performance from the perspective of PT itself with two indicators: the count of successful post-exploit and the time cost for completing 200,000 episodes training of PT. As shown in Fig. 7, after replacing the new and old policies and changing the thread and chief network parameter update mechanism, the post-exploit count of the DPPO model is 276 more than that of DeepExploit, but the time cost only increased by 311 seconds, which indicates that applying DPPO to automate PT is effective and its penetration performance is better than the DeepExploit.

For A3C-GAIL and DPPO-GAIL, we combined the GAIL network then employed the A3C-GAIL to perform intelligent PT. The post-exploit counts of A3C-GAIL are the highest among the four algorithms models. However, the time cost of A3C-GAIL is also the highest, with only 260 times more than the DPPO-GAIL in the count of post-exploit; the time cost is twice as high as DPPO-GAIL, which explains the reason why the total reward value of A3C-GAIL in Fig. 7 is higher than that of DPPO-GAIL. Since we only counted the number of successful post-exploit and gave positive rewards for the exploit that did not establish a session with the target host, A3C-GAIL may have performed multiple exploitations in extra time but failed in the post-exploit. Therefore, considering the count of post-exploit and time cost, we can conclude that DPPO-GAIL performs more efficient intelligent PT.

Time cost under exploiting limited vulnerabilities. In order to verify the better performance of applying the DPPO-GAIL model to perform intelligent PT, we conducted another set of comparative experiments. Since Metasploit can only exploit the vulnerabilities of 11 service types in the target host Metasploitable2, we counted the time it takes to exploit these 11 different vulnerabilities for the four algorithms respectively as the last evaluation indicator. Intuitively, from Fig. 8 we can see that under the guidance of expert samples and a combination of the GAIL network, the time cost of the DPPO-GAIL is the least, followed by A3C-GAIL. What is more, the time cost of DPPO-GAIL, A3C-GAIL and DPPO is reduced by 55.7%, 46.1%, and 17.1% compared with DeepExploit, respectively.

Considering the above five different evaluation indicators adequately, under the guidance of expert sample knowledge and combination with the GAIL network, the performance of A3C-GAIL and DPPO-GAIL in intelligently exploiting the actual target Metasploitable2 has improved in total reward and the
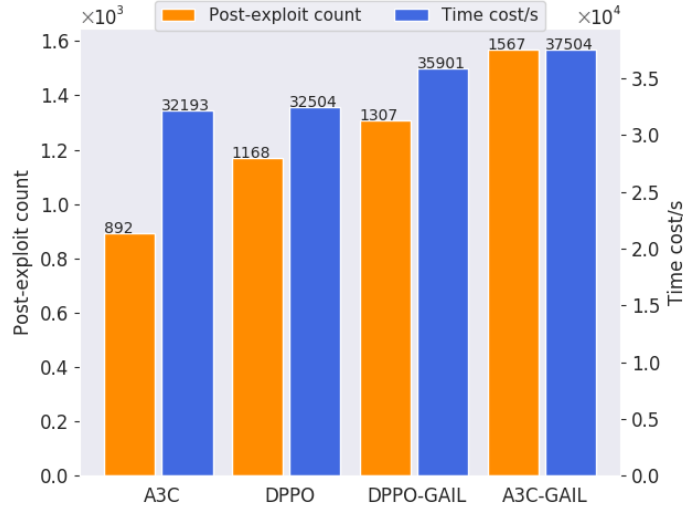
24

Figure 7: Comparison of post-exploit count and time cost of different algorithms in exploiting the target host Metasploitable2

count of successful post-exploit compared with DeepExploit. The total reward value has increased by 253% and 160%, respectively, and the count of post-exploit has increased by 75.7% and 46.5%. However, the time cost of A3C-GAIL was twice as high as that of DPPO-GAIL. Besides, after exploiting 11 different types of service vulnerabilities, the time cost of the two algorithms was reduced by 46.1% and 55.7%, respectively.
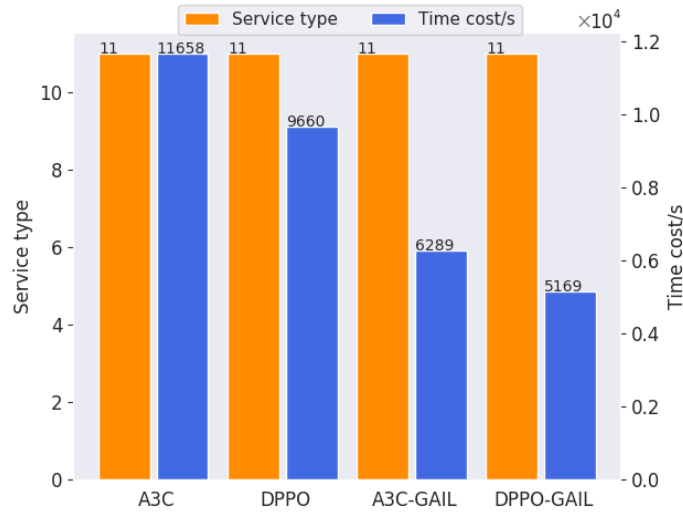
Figure 8: Comparison of and time cost of exploiting 11 different types of service vulnerabilities in the target host Meatsploitable2

Answer to RQ2: The penetration performance of A3C-GAIL and DPPO-GAIL is better than that of A3C and DPPO. Besides, the DPPO-GAIL can execute more post-exploits successfully, and exploit more vulnerabilities of different service types in less time to achieve the SOTA performance.

5.3. Network scenario

In order to verify the effectiveness of automating PT in different scale networks based on GAIL, we conducted another part of experiments on three different network scenarios: a small-scale network with or without honeypot, and a large-scale network in a high-fidelity network simulator Nasim. The primary purpose of this scenario is to generate an optimal penetration path in a trial and error training manner, through exploiting one host in one subnet and gaining control of it, then using the compromised host as a stepping stone according to network connection relationships to carry out several lateral movements for obtaining the control of the final sensitive host.

5.3.1. Experiment settings for different network scenarios

At first, we automate PT in a small-scale network (without honeypot) scenario as shown in Fig. 9. The network contains 4 sub-networks and 8 hosts, sensitive hosts are (2, 0) and (4, 0), which is the target host in PT. Firewalls block communication between subnets with restricted services, and penetrating from one subnet to another requires a specific cost each time.
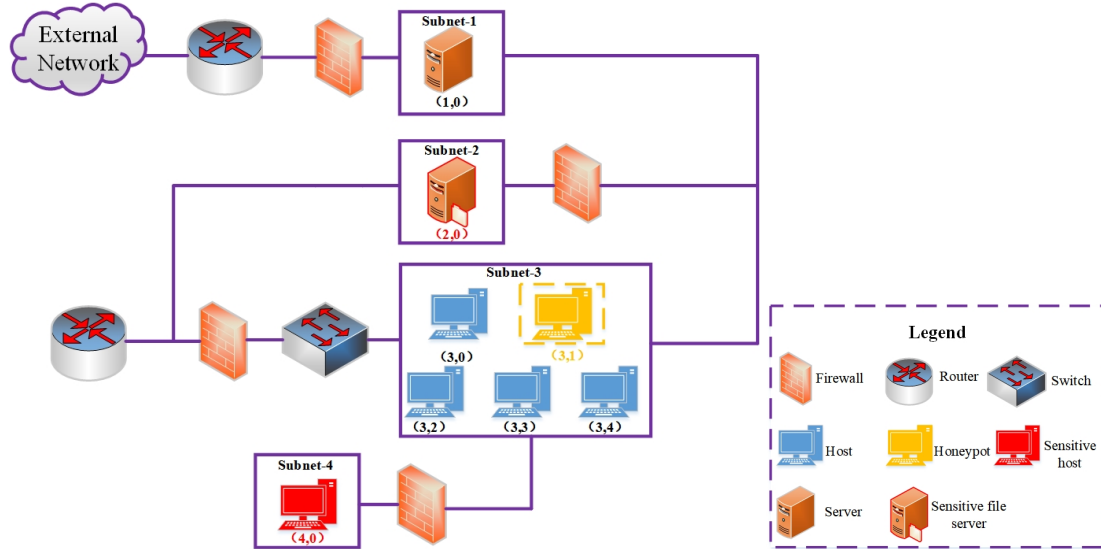
Figure 9: It is the layout of the small-scale network. A honeypot is located in (3,2). In the scene without a honeypot, the location of (3,2) is regarded as a normal host.

The hosts in the subnet can communicate with each other. The configuration information of each host demonstrating in Table 4, including the hosts' respective virtual address, operating system, service information, process information, and the value of the host. In order to simulate the action of an actual attacker, we assume that the agent cannot directly obtain network topology information and host configuration information. Therefore, in addition to processing exploits(Exploit) and privilege promotion (Promotion) actions, scanning (Scan) can also be used to obtain relevant information about the host and target network.

For each host, the agent can choose the actions shown in Table 5. We choose processes and services that attackers often use during the PT process in a subnet to replace network security vulnerabilities, set the probabilities of '0.9' and '0.6' for SSH, HTTP and FTP according to CVSS score, respectively, and set the probability to '1' for privilege promotion and scanning. For example, the agent can get the user permission of the host through exploiting the vulnerability with FTP service then performing the Daclsvc privilege promotion action to get the root permission of the host, thereby realizing a penetration attack on the host.

Table 4: Configuration list

Address	Operation System	Host-value	Service	Process
(1,0)	Linux	0	HTTP	Tomcat
(2,0)	Linux	100	SSH,HTTP	/
(3,0)	Windows	0	HTTP	/
(3,1)	Windows	0	FTP,HTTP	Daclsvc
(3,2)	Windows	0	FTP,HTTP	Daclsvc
(3,3)	Windows	0	FTP	/
(3,4)	Windows	0	FTP	Daclsvc
(4,0)	Linux	100	SSH,HTTP	Tomcat

Then, we conduct the same experiments in the small-scale network with a honeypot scenario shown in Fig. 9 as well, optimizing strategy for finding the penetration path through RL algorithm to realize an intelligent PT process. The honeypot is located in (3,2). But in the samll-scale without a honeypot, the location of (3,2) is regarded as a normal host. We set the value of the honeypot host to -100, assuming the agent knows the address of the honeypot in the subnet. In this case, the agent can bypass the honeypot for exploiting.

Table 5: Agent action list

Name	Type	Operation System	Cost	Prob	Access
SSH-Exp	Exploit	Linux	3	0.9	User
FTP-Exp	Exploit	Windows	1	0.6	User
HTTP-Exp	Exploit	None	2	0.9	User
Tomcat	Promotion	Linux	1	1	Root
Daclsvc	Promotion	Windows	1	1	Root
Service-Scan	Scan	/	1	1	/
Os-Scan	Scan	/	1	1	/
Subnet-Scan	Scan	/	1	1	/
Process-Scan	Scan	/	1	1	/

At last, we conducted the same PT experiment in a large-scale simulated network scenario with 23 hosts. Different from the small-scale network scenario, the large-scale network scenario is more complicated: as shown in Fig. 13, it has 8 subnets and a total of 23 hosts. There is a sensitive host in

subnet 2 and subnet 7, respectively (2, 0) and (7, 0), which are the target hosts of the PT. The hosts in the remaining subnets are all normal hosts; each host runs various services and communicates. Besides, only subnet 1 could communicate with the external network directly. The remaining 6 subnets communicated in the internal network; subnet 5 could only communicate with subnet 3 directly, and subnet 6 could only communicate with subnet 4 directly. For example, if the attacker arrives at (5,0), his goal is to exploit the sensitive host (7,0). Therefore, he has to come back to subnet 3 for lateral movements and other operations to exploit one host in subnet 4, then arrive at (7.0) through the same operation. It significantly increases the cost of the attack, so our purpose is to find the optimal penetration path at the lowest cost.

Besides, there are three operating systems types: Windows, Linux and Unix; seven different exploiting services: SSH, FTP, HTTP, RPC, PHP, Samba, SMTP (Simple Mail Transfer Protocol), and three privilege promotion process: Tomcat, Daclsvc, and Schtask in the large-scale network scenario. We set the probabilities of '0.9' for SSH, HTTP; set '0.6' for FTP, RPC, PHP and SMTP; set '0.3' for Samba according to CVSS score respectively, and set the probability to '1' for privilege promotion and scanning. The penetration mechanism of this scene is similar to the small-scale network.

However, due to the increase of exploiting services and hosts, the input state in the large-scale scene is expanded to 768 dimensions, and also the output action is expanded to 322 dimensions, which is a high-dimensional input and output scene. Due to the higher complexity of large-scale network scenarios, the combination of reinforcement learning to optimize the penetration path has the problem of too large state dimensions and sparse reward values. Therefore, we also introduce the expert sample knowledge into this scene and combine the GAIL network with training the agent for finding and optimizing the penetration path.
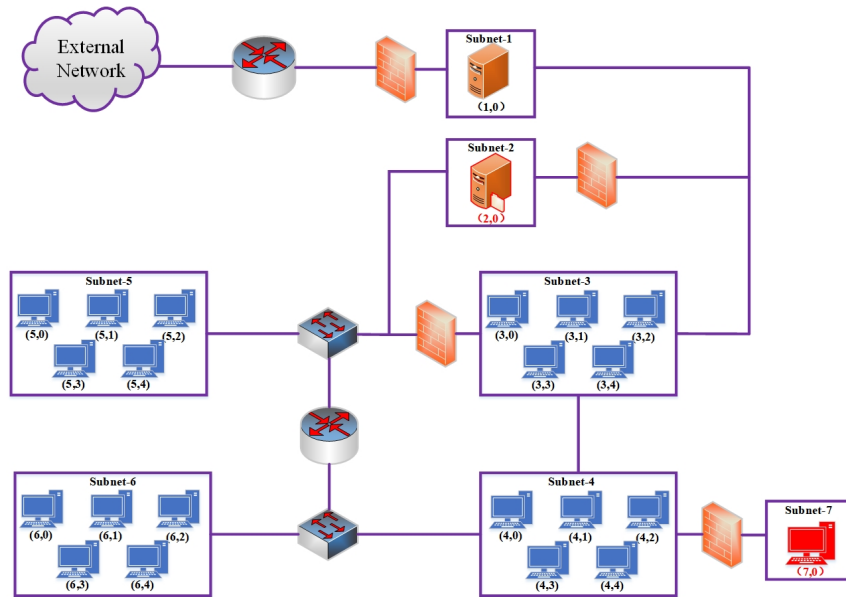

Figure 10: It is the layout of the large-scale network. Sensitive file server and sensitive host located in (2,0) and (7,0), respectively.

The reward value of successfully exploiting the sensitive host is 100 in the small scale network with or without the honeypot scenario illustrated in Fig. 9, but the reward value of falling into the honeypot is set to -100. Likewise, in the large-scale network shown in Fig. 13, setting the reward values of successfully exploiting each of two sensitive hosts to 100. Agents can only execute exploit between connected subnets or hosts in the network. The ending of each round of the agent are the following two situations:

- Obtaining root permissions for all sensitive hosts.

- Training steps in each round have reached the maximum value.

The agent is tried to obtain the maximum cumulative reward value through continuous trial and error, thereby learning the optimal penetration path for exploiting the target sensitive host. In the small-scale with a honeypot network, since the reward value of falling into the honeypot is negative, the agent will choose to bypass the honeypot during the PT process and finally learn the optimal penetration path to exploit the sensitive host with fewer steps.

We first use the Q-learning algorithm to conduct experiments and then perform automated PT on the small-scale, small-scale with honeypot and large-scale networks. Second, expert knowledge is introduced based on the benchmark algorithm Q-learning and combines the GAIL network with Q-learning. At last, compare the penetration performance of Q-learning-GAIL and Q-learning in finding the optimal penetration path. Besides, we set the training round is 20000 for each scenario, recorded the steps and cost needed to generate the penetration path in each round, and adopted the average value as the experiment results. For two RL algorithms in three different scenarios: Q-learning and Q-learning-GAIL, we apply them to automate PT and evaluate their performance.

5.3.2. Evaluation metrics

In the training phase, we calculating the average value every 1000 rounds, and applying the following three indicators to analyze the answers of ***RQ3*** and ***RQ4***:

- **Average reward.** In the training process, the average round reward is calculated every certain round, which can be used to evaluate the overall performance of the algorithm;

- **Average steps.** In the training process, the average round length is calculated every certain round, which represents the time cost required for the generation of the penetration path, and is also used to evaluate the overall performance of the algorithm.

- **Probability of invading honeypot.** In the process of PT, the average probability of honeypot intrusion is calculated every certain round to evaluate the effectiveness of honeypot deployment.

5.3.3. Research question 3

Does expanding the expert sample amount still improve the penetration performance in different network scenarios?

We will answer the question through analyzing the experiment results from the one indicator: Training rounds.

The optimal path is definite when optimizing the penetration path in a network scene. Therefore, there is no situation where a particular state corresponds to multiple optimal penetration paths. In other words, when the agent finds the best penetration path, the state and action at this time are a one-to-one relationship. What is more, there is no concept of the complexity of expert knowledge rules. Therefore, we are only considering the impact of expert sample amount on penetration performance.

We evaluated the impact of different expert sample amounts on the improvement of penetration performance in three different network scenarios separately and used the training rounds required to find the optimal penetration path as the indicator. It can be seen from Table 6 that the sample size of expert knowledge has a specific impact on the training effect of optimizing the penetration path. Without introducing expert sample knowledge, the training rounds required to find the optimal penetration path in the small-scale network, small-scale with honeypot network scenario, and large-scale network scenario are 4000, 5000, and 3000, respectively. When the sample amount expands to 5000 pairs, the rounds for generating the optimal penetration path are 3000, 3500, and 2000, respectively, reduced by 1000, 1500, and 1000 rounds in turn.

On this basis, the expansion of the expert knowledge sample size did not significantly accelerate generating the optimal penetration path. Therefore, we all introduced 5000 pairs of expert knowledge sample amounts in the following three scene experiments.

Table 6: Training rounds for different expert sample amount to find the
optimal penetration path in different network scenarios

Network scene	Expert sample amount 0	2500	5000	7500
Small-scale	4000	3500	**3000**	3000
Small-scale with honeypot	5000	4000	**3500**	3500
Large-scale	3000	3000	**2000**	2000

Answer to RQ3: Penetration performance improved with expanding the expert samples in the range of 5,000 pairs, but it is not significantly improved when the expert sample exceeds this range.
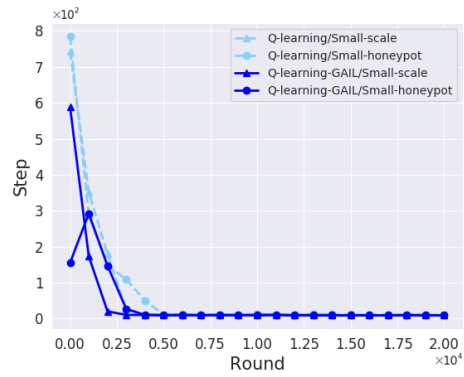
5.3.4. Research question 4

Can GAIL also achieve the SOTA penetration performance in different network scenarios?

We will answer this question from three different network scenarios combined with the three indicators: Average reward, Average step and Probability of invading honeypot through the following experimental results.
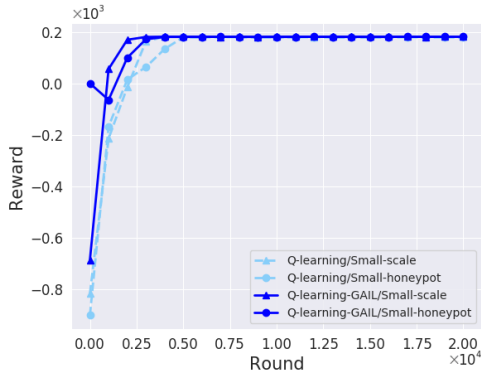
(a) Comparison of average reward   (b) Comparison of average steps

Figure 11: Comparison of average reward (left) and average steps (right) in small-scale network with or without honeypot scenarios by Q-learning and Q-learning-GAIL with 20000 rounds training.

Small-scale network It can be seen from the comparison of the average reward value and average steps in Fig. 11, with the introduction of expert knowledge samples and the combination of the GAIL mechanism, accelerate the convergence of the reward function and training steps, and finally improve the penetration performance to a certain extent. In a small-scale network scenario without a honeypot, using Q-learning-GAIL algorithm to optimize the penetration path, the average reward value began to converge in 2000 rounds, and the average training steps also showed the same trend, both of them are about 1000 rounds faster than the Q-learning.

As to small-scale with a honeypot, the average reward value of each round is higher than Q-learning. The fluctuation of the reward value before convergence is much gentle than that of Q-learning. In terms of training steps, the average round steps of the Q-learning-GAIL algorithm tend to converge at 3000 rounds, while Q-learning tends to converge at 5000 rounds. The average reward value also shows this trend, and both are about 2000 rounds faster than Q-learning for finding the optimal penetration path.

In short, the introduction of expert knowledge samples speeds up finding the optimal penetration path in the small-scale network. Because the complexity of the small network is relatively lower than large-scale, the introduction of the GAIL mechanism does not reduce the average round steps and the steps of finding optimal penetration stabilizing at 9 steps in the end. On the other side, training costs are also reduced. Especially in the honeypot scenario, the reduction in training cost is twice that in the case of no honeypot.

Small-scale network with honeypot In addition, this section comparative analyzes the intrusion probability of honeypots in the small-scale network with a honeypot. Intuitively, this indicator reflects the effectiveness of deploying honeypots, that is, whether the agent will fall into the honeypot again after training through expert knowledge and GAIL network guidance. The agent's learning process aims to exploit sensitive hosts, not to invade the honeypot. We set invading the honeypot (3, 2) a negative reward. Therefore, the lower the probability of falling into the honeypot, the better the expert sample guides the
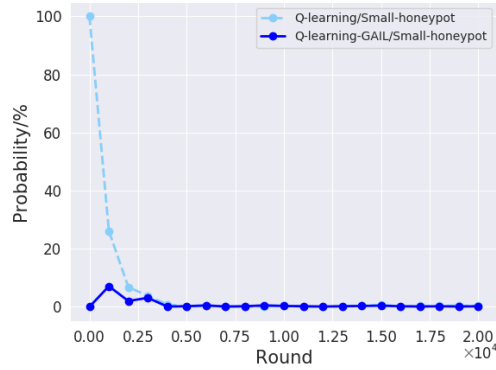
Figure 12: Probability comparison of invading the honeypot

agent to train, which also indicates that the agent's action distribution is infinitely close to the expert sample at this time.

Fig. 12 intuitively reflects that the probability that the agent trained by Q-learning-GAIL will fall into the honeypot is much lower than the benchmark Q-learning in the first 3000 rounds. At the beginning of training, the probability of falling into the honeypot is 100% training by Q-learning, while training with Q-learning-GAIL is directly reduced to 0. Although in the next 3000 rounds, the probability that GAIL guides the agent to fall into a honeypot fluctuates slightly, in the entire 20000 rounds, the probability of applying Q-learning-GAIL algorithm to fall into honeypot is lower than Q-learning in each round, which also indicates that the introduction of the GAIL mechanism improves the penetration performance more significantly in the network scenarios with honeypot.

Large-scale network Compared with the small-scale network, it is not difficult to find that, when applying Q-learning-GAIL to determine optimal penetration path, the improvement of PT performance is more evident in the large-scale network, as shown in Fig. 13.

In the early stage of training, the agent is still in the exploring stage. The average reward and average steps of the benchmark Q-learning in the first 3000 rounds show a major fluctuations trend. However, under the guidance of expert knowledge, the Q-learning-GAIL algorithm shows a relatively stable performance in the early stage of the training agent for optimal penetration path-finding. The average reward value is positive in the first 3000 rounds, and the average exploration steps are within 50 steps. After 2000 rounds, both the average reward value and the average steps show a convergence trend. Moreover, the final average steps stabilized at 15 steps, while the final average steps of Q-learning is 17-18 steps, which 2-3 steps have accelerated in the final round. In short, by introducing expert knowledge and combining the GAIL network with Q-learning to optimize penetration path, Q-learning-GAIL shows a more stable penetration performance and achieves the SOTA penetration performance with less time and cost.

Answer to RQ4: Whether in a small-scale network with honeypot or not or in a large-scale network, the penetration performance of Q-learning-GAIL has all reached the SOTA. Especially in small-scale with honeypot and large-scale networks, penetration performance improved even more apparent.
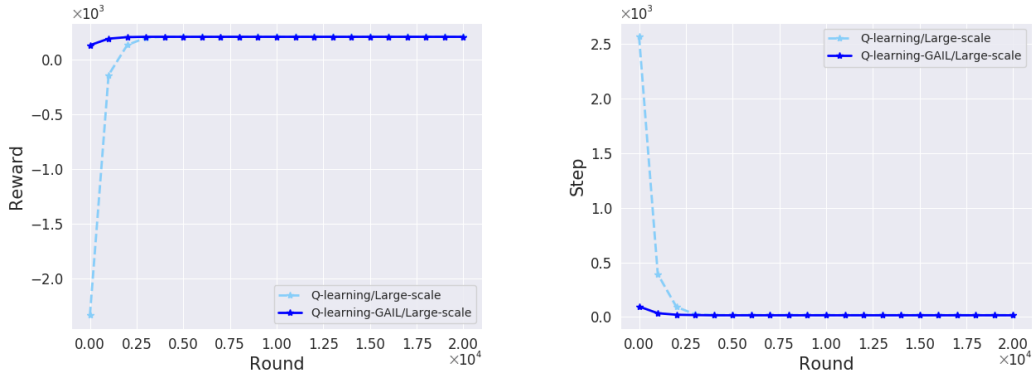
(a) Comparison of average reward

(b) Comparison of average steps

Figure 13: Comparison of average reward (left) and average steps (right) in large-scale network scenarios

6. Conclusion and future work

This paper proposed a method to automate the penetration testing process based on GAIL, and evaluated the effectiveness and performance of exploiting the single host and three different networks. The imitation learning agent is modeled as a penetration attacker in the actual and simulated scenarios. In the experiments, we compared the methods of A3C-GAIL, DPPO-GAIL and Q-learning-GAIL, which are our methods, with the DeepExploit, DPPO and Q-learning in exploiting single host and network. The results testified the performance of our methods.In addition to combining with A3C, DPPO and Q-learning, GAIL can also combine with other DRL / RL algorithms. What is more, our original intention is to propose a general penetration testing framework GAIL-PT that can help conjunct GAIL with different DRL / RL algorithms to automate the penetration testing process in different scenarios. To our best knowledge, it is the first study of applying GAIL for penetration testing. We publish the code on Github so that interested scholars can conduct related research.

In aspect of limitation of the proposed method, its complexity of A3C-GAIL and DPPO-GAIL is relatively high. Although the network simulator NASim appears accurate vulnerability service and specific penetration operations in network scenarios, there are still differences between NASim and the existing network, and our method has not yet been validated in real complex networks. Therefore, future work can be carried out in the following two aspects: reducing the algorithm's complexity and applying GAIL-PT to exploit the real complex network.

7. Acknowledgment

References

[1] Arce, I., McGraw, G., 2004. Guest editors' introduction: Why attacking systems is a good idea. IEEE Security & Privacy 2, 17–19.

[2] Arkin, B., Stender, S., McGraw, G., 2005. Software penetration testing. IEEE Security Privacy 3, 84–87. doi:`10.1109/MSP.2005.23`.

[3] Arora, S., Doshi, P., 2021. A survey of inverse reinforcement learning: Challenges, methods and progress. Artificial Intelligence , 103500.

[4] Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al., 2019. Dota 2 with large scale deep reinforcement learning. arXiv preprint arXiv:1912.06680 .

[5] Bland, J.A., Petty, M.D., Whitaker, T.S., Maxwell, K.P., Cantrell, W.A., 2020. Machine learning cyberattack and defense strategies. Computers & security 92, 101738.

[6] Canese, L., Cardarilli, G.C., Di Nunzio, L., Fazzolari, R., Giardino, D., Re, M., Spanò, S., 2021. Multi-agent reinforcement learning: A review of challenges and applications. Applied Sciences 11, 4948.

[7] Dulac-Arnold, G., Evans, R., van Hasselt, H., Sunehag, P., Lillicrap, T., Hunt, J., Mann, T., Weber, T., Degris, T., Coppin, B., 2015. Deep reinforcement learning in large discrete action spaces. arXiv preprint arXiv:1512.07679 .

[8] El Kamel, N., Eddabbah, M., Lmoumen, Y., Touahni, R., 2020. A smart agent design for cyber security based on honeypot and machine learning. Security and Communication Networks 2020.

[9] Elderman, R., Pater, L.J., Thie, A.S., Drugan, M.M., Wiering, M.A., 2017. Adversarial reinforcement learning in a cyber security simulation., in: ICAART (2), pp. 559–566.

[10] Fan, J., Wang, Z., Xie, Y., Yang, Z., 2020. A theoretical analysis of deep q-learning, in: Learning for Dynamics and Control, PMLR. pp. 486–489.

[11] Farquhar, G., Gustafson, L., Lin, Z., Whiteson, S., Usunier, N., Synnaeve, G., 2020. Growing action spaces, in: International Conference on Machine Learning, PMLR. pp. 3040–3051.

[12] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. Advances in neural information processing systems 27.

[13] Haeni, R.E., 1997. Firewall penetration testing. Technical Report. Citeseer.

[14] He, X., Dai, H., Ning, P., 2016. Faster learning and adaptation in security games by exploiting information asymmetry. IEEE Transactions on Signal Processing 64, 3429–3443.

[15] HelpSysthems, 2021. Core impact. `https://www.coresecurity.com/products/core-impact/`.

[16] Ho, J., Ermon, S., 2016. Generative adversarial imitation learning. Advances in neural information processing systems 29, 4565–4573.

[17] Ho, J., Gupta, J.K., Ermon, S., 2016. Model-Free Imitation Learning with Policy Optimization. arXiv e-prints , arXiv:1605.08478`arXiv:1605.08478`.

[18] Hu, Z., Beuran, R., Tan, Y., 2020. Automated penetration testing using deep reinforcement learning, in: 2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), IEEE. pp. 2–10.

[19] Ingols, K., Chu, M., Lippmann, R., Webster, S., Boyer, S., 2009. Modeling modern network attacks and countermeasures using attack graphs, in: 2009 Annual Computer Security Applications Conference, IEEE. pp. 117–126.

[20] Isao Takaesu, 2018. Metasploit meets machine learning. `https://www.mbsd.jp/blog/20180228.html`.

[21] Kaur, G., Kaur, N., 2017. Penetration testing–reconnaissance with nmap tool. International Journal of Advanced Research in Computer Science 8.

[22] Kaushik, M., Ojha, G., 2014. Attack penetration system for sql injection. International journal of advanced computer research 4, 724.

[23] Kennedy, D., O'gorman, J., Kearns, D., Aharoni, M., 2011. Metasploit: the penetration tester's guide. No Starch Press.

[24] Levine, S., Popovic, Z., Koltun, V., 2011. Nonlinear inverse reinforcement learning with gaussian processes. Advances in neural information processing systems 24, 19–27.

[25] Mapper, N., 2021. Nmap. `https://nmap.org/`.

[26] McDaniel, L., Talvi, E., Hay, B., 2016. Capture the flag as cyber security introduction, in: 2016 49th hawaii international conference on system sciences (hicss), IEEE. pp. 5479–5486.

[27] Mell, P., Scarfone, K., Romanosky, S., et al., 2007. A complete guide to the common vulnerability scoring system version 2.0, in: Published by FIRST-forum of incident response and security teams, p. 23.

[28] Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K., 2016. Asynchronous methods for deep reinforcement learning, in: International conference on machine learning, PMLR. pp. 1928–1937.

[29] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M., 2013. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602 .

[30] Moyer, P.R., Schultz, E.E., 1996. A systematic methodology for firewall penetration testing. Network Security 1996, 11–18.

[31] Neal, C., Dagdougui, H., Lodi, A., Fernandez, J.M., 2021. Reinforcement learning based penetration testing of a microgrid control algorithm, in: 2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC), IEEE. pp. 0038–0044.

[32] Phillips, C., Swiler, L.P., 1998. A graph-based system for network-vulnerability analysis, in: Proceedings of the 1998 workshop on New security paradigms, pp. 71–79.

[33] Pozdniakov, K., Alonso, E., Stankovic, V., Tam, K., Jones, K., 2020. Smart security audit: Reinforcement learning with a deep neural network approximator, in: 2020 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA), IEEE. pp. 1–8.

[34] Qiu, X., Jia, Q., Wang, S., Xia, C., Lv, L., 2014a. Automatic generation algorithm of penetration graph in penetration testing, in: 2014 Ninth International Conference on P2P, Parallel, Grid, Cloud and Internet Computing, IEEE. pp. 531–537.

[35] Qiu, X., Wang, S., Jia, Q., Xia, C., Xia, Q., 2014b. An automated method of penetration testing, in: 2014 IEEE Computers, Communications and IT Applications Conference, IEEE. pp. 211–216.

[36] RAPID7, 2017. Metasploitable2. `https://docs.rapid7.com/metasploit/metasploitable-2/Download:https://docs.rapid7.com/metasploit/metasploitable-2/`.

[37] RAPID7, 2019. Nexpose. `https://www.rapid7.com/products/nexpose/`.

[38] RAPID7, 2021. Metasploit. `http://www.metasploit.com/`.

[39] Sarraute, C., Buffet, O., Hoffmann, J., 2012. Pomdps make better hackers: Accounting for uncertainty in penetration testing, in: Twenty-Sixth AAAI Conference on Artificial Intelligence.

[40] Sarraute, C., Buffet, O., Hoffmann, J., 2013. Penetration testing== pomdp solving? arXiv preprint arXiv:1306.4714 .

[41] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 .

[42] Schwartz, J., Kurniawati, H., 2019. Autonomous penetration testing using reinforcement learning. arXiv preprint arXiv:1905.05965 .

[43] Schwartz, J., Kurniawati, H., El-Mahassni, E., 2020. Pomdp+ information-decay: Incorporating defender's behaviour in autonomous penetration testing, in: Proceedings of the International Conference on Automated Planning and Scheduling, pp. 235–243.

[44] Schwartz, J., Kurniawatti, H., 2019. Nasim: Network attack simulator. `https://networkattacksimulator.readthedocs.io/`.

[45] Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al., 2016a. Mastering the game of go with deep neural networks and tree search. nature 529, 484–489.

[46] Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al., 2016b. Mastering the game of go with deep neural networks and tree search. nature 529, 484–489.

[47] Spitzner, L., 2003. Honeypots: tracking hackers. volume 1. Addison-Wesley Reading.

[48] Sun, R., Li, X., Zhang, L., Liu, J., 2020. Distributed storage codes based on double-layered piggy-backing framework. IEEE Access 8, 150447–150464.

[49] Tenale, 2021. Nessus. `https://zh-cn.tenable.com/products/nessus`.

[50] Torabi, F., Warnell, G., Stone, P., 2018. Behavioral cloning from observation. arXiv preprint arXiv:1805.01954 .

[51] Tran, K., Akella, A., Standen, M., Kim, J., Bowman, D., Richer, T., Lin, C.T., 2021. Deep hierarchical reinforcement agents for automated penetration testing. arXiv preprint arXiv:2109.06449 .

[52] Vinyals, O., Babuschkin, I., Czarnecki, W.M., Mathieu, M., Dudzik, A., Chung, J., Choi, D.H., Powell, R., Ewalds, T., Georgiev, P., et al., 2019. Grandmaster level in starcraft ii using multi-agent reinforcement learning. Nature 575, 350–354.

[53] Wang, R., Ciliberto, C., Amadori, P.V., Demiris, Y., 2019. Random expert distillation: Imitation learning via expert policy support estimation, in: International Conference on Machine Learning, PMLR. pp. 6536–6544.

[54] Wiering, M.A., Van Otterlo, M., 2012. Reinforcement learning. Adaptation, learning, and optimization 12.

[55] Zennaro, F.M., Erdodi, L., 2020. Modeling penetration testing with reinforcement learning using capture-the-flag challenges and tabular q-learning. arXiv preprint arXiv:2005.12632 .

[56] ZHANG, Y., Tianyang, Z., Junhu, Z., Qingxian, W., 2020. Domain-independent intelligent planning technology and its application to automated penetration testing oriented attack path discovery. Journal of Electronics & Information Technology 42, 2095–2107.

[57] Zhou, S., Liu, J., Hou, D., Zhong, X., Zhang, Y., 2021. Autonomous penetration testing based on improved deep q-network. Applied Sciences 11, 8823.

[58] Ziebart, B.D., Maas, A.L., Bagnell, J.A., Dey, A.K., et al., 2008. Maximum entropy inverse reinforcement learning., in: Aaai, Chicago, IL, USA. pp. 1433–1438.