

Inroads into Autonomous Network Defence using Explained Reinforcement Learning

Myles Foley¹, Mia Wang¹, Zoe M², Chris Hicks² and Vasilios Mavroudis²

¹Imperial College London

²The Alan Turing Institute

Abstract

Computer network defence is a complicated task that has necessitated a high degree of human involvement. However, with recent advancements in machine learning, fully autonomous network defence is becoming increasingly plausible. This paper introduces an end-to-end methodology for studying attack strategies, designing defence agents and explaining their operation. First, using state diagrams, we visualise adversarial behaviour to gain insight about potential points of intervention and inform the design of our defensive models. We opt to use a set of deep reinforcement learning agents trained on different parts of the task and organised in a shallow hierarchy. Our evaluation shows that the resulting design achieves a substantial performance improvement compared to prior work. Finally, to better investigate the decision-making process of our agents, we complete our analysis with a feature ablation and importance study.

Keywords

Reinforcement Learning, Autonomous Cyber Defence, Deep Learning, Network Defence

1. Introduction

Computer network security is characterised by an asymmetry as the defender needs to ensure constant protection of the network's components, while the adversary can opportunistically single-out weak entry points. Such asymmetries have been identified and addressed in many other areas of cyber security. For example, cryptographic protocols (e.g., TLS) thwart denial of service attacks by ensuring that the prover commits enough computation cycles before the verifier does so. In network defence, however, the problem remains open as the task is complex [1] and involves a wide array of both attack vectors and mitigation tools. Thus, network defence is currently handled primarily by human experts which entails high operational costs.

RL, and particularly deep RL (DRL), excels in interactive tasks that cannot easily be solved using analytical solutions. Human and even super-human levels of performance have been achieved in a range of complex tasks including classic board games such as chess and Go [2, 3], video games ranging from classic Atari [4, 5] to multi-player real-time strategy games [6], autonomous driving [7], and robotics [8]. Recently, DRL has also been successfully applied to

CAMLIS'22: *Conference on Applied Machine Learning in Information Security (CAMLIS)*, October 20–21, 2022, Arlington, VA

✉ m.foley20@imperial.ac.uk (M. Foley); yixuan.wang18@imperial.ac.uk (M. Wang); zm@turing.ac.uk (Z. M); c.hicks@turing.ac.uk (C. Hicks); c.hicks@turing.ac.uk (V. Mavroudis)

ORCID 0000-0002-0877-7063 (M. Foley); 0000-0002-0877-7063 (M. Wang)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

autonomous network defence [9], a highly interactive task where the defender proactively monitors the state of the network, identifies abnormalities, and acts to remediate them. Commonly, this takes the form of a shallow hierarchy of *specialised subagents* coordinated by a *controller*, any combination of these being autonomous.

To date however, there has been limited consideration for the explainability of these models. Explainable AI has, in domains such as natural language processing and computer vision [10], proven useful not only for end users but also experts and developers of AI systems. DRL models are particularly challenging to explain because the neural networks which represent their agent policies are not readily understandable by humans. Nonetheless, the ability to explain and understand the actions of an autonomous defensive agent is critical. This work investigates, and answers in the affirmative, whether explainable RL (XRL) models and environments can improve autonomous defensive capabilities and aid in their development.

1.0.1. Contributions

Our main contributions are:

- We develop methodologies for visualising (i.e., explaining) attacker functionality in the CybORG cyber environment. Our methodology highlights previously undocumented differences in the adversary models and motivates two new controller architectures with improved classification accuracy.
- We present the full details of our new controller and specialised subagent models. We then evaluate them against two classes of adversary in the CybORG environment realising substantial performance improvements.
- We perform a feature ablation and importance study to understand the most influential elements in the observation space and explain our model outputs.

2. RL Background

In this section we discuss the key RL techniques that are relevant for the rest of the paper.

2.1. Deep RL Algorithms

2.1.1. PPO

Proximal Policy Optimisation (PPO) is an efficient policy gradient method [5] for DRL. It has been shown to outperform other popular algorithms such as A3C [2], achieving super-human performance in a variety of complex tasks including 49 separate ATARI arcade games [5]. Despite its effectiveness in very complex environments [11], it has seen only limited use in security settings [12, 13],

PPO uses a policy π_θ ($\theta \in \mathbb{R}$) with an objective function that is defined by the total reward $J(\theta) = \mathbb{E}_{\pi_0}[\sum_{t=0}^{\infty} \gamma^t r_t]$. By formulating the objective function in this way actor-critic architectures can be used: the actor selects an action which is evaluated by the critic. The policy gradient is then computed:

$$\nabla_\theta J(\theta) = \mathbb{E}_{\pi_0}[\nabla_\theta \log \pi_\theta(s, a) A_{\pi_\theta}(s)] \quad (1)$$

$A_{\pi_\theta}(s)$ is the advantage of taking action a instead of the *average* action as computed by the policy π_θ ($Q_{\pi_\theta}(s, a) - V_{\pi_\theta}(s)$) [14].

During gradient descent, PPO introduces a clipping function to both prevent reaching local optima during large updates and avoid smaller updates that significantly increase the length of training.

2.2. Curious Exploration

Curiosity is a technique that enables agents to explore their environment based on an intrinsic reward signal not provided by the environment [15]. Such a signal is particularly useful in the absence of a continual extrinsic reward (e.g., the running score found in some games). Pathak et al. [15] introduce the Intrinsic Curiosity Module (ICM), a semi-supervised technique in which agents choose actions based on the uncertainty in the outcome of each action, intrinsically motivating the exploration of unknown states. ICM also ensures that agents are only incentivised to reach states that are impacted by their actions, avoiding those which are inherently unpredictable.

2.3. Explainable RL

Explainable RL (XRL), a fledgling sub-field of explainable AI, is the study of tools and methods which enhance human understanding of the actions taken by autonomous agents. A recent and thorough review of XRL is provided by Heuillet et al. [16] and separately by Puiutta and Veith [17]. XRL methods are commonly divided between those which are intrinsic, sometimes called transparent, and those which are post-hoc. Intrinsic XRL models are inherently interpretable and offer explainability at the time of training. In contrast, post-hoc explainability occurs after training; often by creating a second, simpler model to provide explanations. In DRL, learned policies are represented by neural networks making them difficult to interpret. Post-hoc explainability allows the performance advantages of DRL [3] to be retained whilst facilitating human understanding of autonomous decision making. Explainability is not limited to users and experts affected by the decisions of models but, as in this work, is a valuable researcher's aid in developing more efficient and higher-performance models.

3. Network Simulation Environment

We use the CybORG environment [18] which simulates the computer network of a manufacturing plant, as shown in Figure 1. The network consists of five *user hosts* (Subnet 1), three *enterprise servers* (Subnet 2¹), three *operational hosts* and the *operational server* (Subnet 3). Each host exposes a number of network services that other hosts can connect to, and which may have exploitable vulnerabilities. However, due to the network's firewalls hosts in Subnet 1 cannot directly connect to machines in Subnet 3, and the operational server is accessible only through the operational hosts. The liveness of the operational server has a direct impact on the manufacturing and is considered critical. CybORG assumes two players, a defender and an adversary, who interact with the turn based environment using the actions available to them.

¹Subnet 2 also includes the defender's machine.

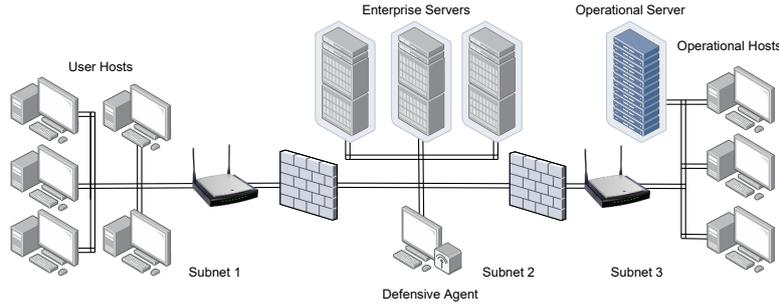


Figure 1: The CybORG environment showing the three subnets and their corresponding hosts and firewalls.

A common drawback of simulated environments in RL is the *reality gap* which causes agents not to generalise sufficiently when moved from the simulation (i.e., training) to reality (i.e., evaluation). This is due to the simulation not adequately matching reality (e.g., in robotics). To address this, CybORG provides a network *emulator* that runs on Amazon Web Services (AWS). The combination of simulation and emulation ensure that the reality gap is minimised, with the actions available and their effect on the environment consistent across both [18].

The CybORG environment is host to the ‘Cyber Autonomy Gym for Experimentation’ (CAGE) challenge [19, 20, 21]. CAGE is an international Kaggle-style competition, providing an increasingly challenging benchmark for the evaluation of autonomous defensive agents. The competition is currently in its second iteration (CAGE II).

3.1. Action Space

Attackers and defenders have unique action spaces. Defenders perform actions at the host level: 1) Analysing the processes running, 2) Terminating malicious processes, 3) Restoring the host to a previous (benign) state, and 4) Deploying *honeypot*² services. Adversaries can: 1) Scan a subnet for hosts, 2) Scan the ports of a host, 3) Exploit a service on a port, 4) Escalate their access, and 5) Disrupt the services on the operational server. Both players have a ‘sleep’ action to perform no action on the network. Based on the selected actions, the environment updates its *state* and updates the agents’ scores. It should be noted that even valid actions may not succeed, as the CybORG simulator introduces randomness to mimic the behaviour of the emulator (e.g. a valid node restoration may occasionally fail).

3.2. Observation Space

The defender’s observation space is a vector of 52 bits i.e., 4 bits for every network host. The first two bits represent whether the host state is unknown (none), scanned or exploited (set when a decoy is triggered); the last 2 bits specify the access the attacker has on the host machine (i.e., none, user and administrator).

²Honeypot refers to a decoy system or service that lures attackers by appearing to suffer from known security vulnerabilities. Honeypots are used to detect malicious actors and study their behaviour.

As in a real network defence situation, neither the defender nor the adversary is omniscient. Neither agent knows the state of the network or the other's position with absolute certainty. In addition, the outcomes of actions are stochastic. For example, from the defender's perspective, when an exploit fails it is not possible to precisely determine which exploit was attempted. This can be crucial information in the instance that an adversary favours a specific exploit strategy. A better informed defender could strategically place decoys on the targeted service to frustrate and evade further attempts more effectively.

3.3. Reward Function

Most games include a scoring function that quantifies the performance of the player. Similarly, Cyborg uses a reward function that rewards the adversary and penalises the defender for every compromised or impacted network host. The reward function is as follows: on each turn, for every host on which the adversary has admin access, the defender receives a reward of -0.1 and for every server the reward is -1. There is a -10 reward for disruption on the operational server and a -1 reward when any device is 'restored'. In the context of RL, the negative reward for the defensive agent incentivises the agent to take actions that minimise the effect of the adversary.

3.4. Adversaries

The environment includes two adversaries: the `BLineAgent` that has prior knowledge (i.e., full knowledge of the network's structure but not its current state), and the `MeanderAgent` which does not have any prior information. Both agents share the same objective, to reach the operational server and, after escalating their privileges, disrupt its services (i.e., impact its liveness). Due to prior knowledge, the `BLineAgent` follows an optimal exploitation trajectory to the operational server. In contrast, the `MeanderAgent` breadth-wise scans the network for vulnerable hosts and gradually traverses the subnets. To prevent trivial defence strategies, the adversary is given user access on a predetermined host (in Subnet 1) that cannot be 'restored' to a benign state by the defender.

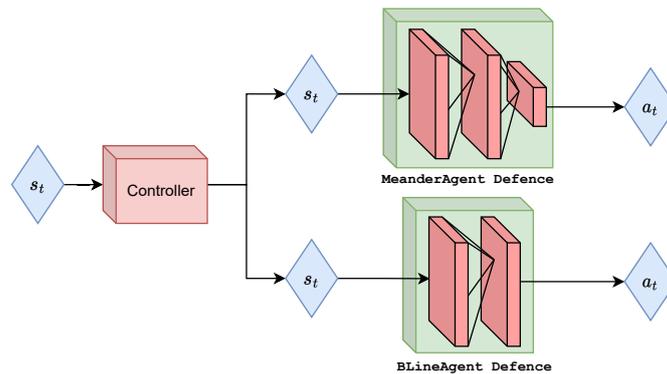


Figure 2: Hierarchical structure of the overall defensive model including the specialist subagents.

4. Model

The models that we train have a similar basic structure to those described in [9] that were trained for CAGE I. In particular, we focus our efforts on training a hierarchy of specialised defensive agents using DRL. These agents feature a controller agent that, at each time step, chooses a subagent to perform the action. Each subagent is trained against a specific adversarial strategy.

As described in Section 3.4, the environment includes two adversaries. The hierarchical architecture was developed specifically to exploit this. The model supports two expert subagents that, through the controller, are ‘consulted’ over the course of an episode (Figure 2). This avoids the performance limitations of a single, more general agent. Given the differences in the two adversaries, each subagent requires a different neural architecture for best performance. These are described below.

4.1. MeanderAgent Defence

Our `MeanderAgent` defensive subagent was trained using the PPO algorithm and utilises a comparatively deeper neural network including three hidden layers with widths 256, 256, and 52. Full details of the hyperparameters used can be found in Appendix A.

Notably, curiosity did not improve the performance. Since the `MeanderAgent` is explicitly designed to explore the network during its attack, the opposing defender is also forced to explore more broadly and to employ a wider range of strategies during training. As such, it learns sufficiently general strategies without the need for curiosity.

4.2. BLineAgent Defence

In contrast to the `MeanderAgent`, the `BLineAgent` follows a near-optimal path through the network. The `BLineAgent` defence, therefore, is at much greater risk of overfitting during training. As a result, we found that when training defensive agents against the `BLineAgent`, it was beneficial to include the curiosity mechanism. In this paper we consider two subagents for `BLineAgent` defence: an Action Knowledge (AK) subagent, and a State Representation (SR) subagent. Both are trained using PPO with curiosity but make different modifications to the state space.

The AK subagent modifies each observation by appending a single bit indicating the success of the previous action. We find that this gives the subagent a better understanding of the defensive process and results in an improvement in performance.

Secondly the SR agent is identical to the AK subagent, but receives observations of 27 floats as opposed to 53 bits. In this state space, each host has two floats to represent the features of activity and compromise. The additional float indicates whether the previous action succeeded. Although the mean episode reward is comparable to the AK agent’s mean reward, we see a notable decrease in variance.

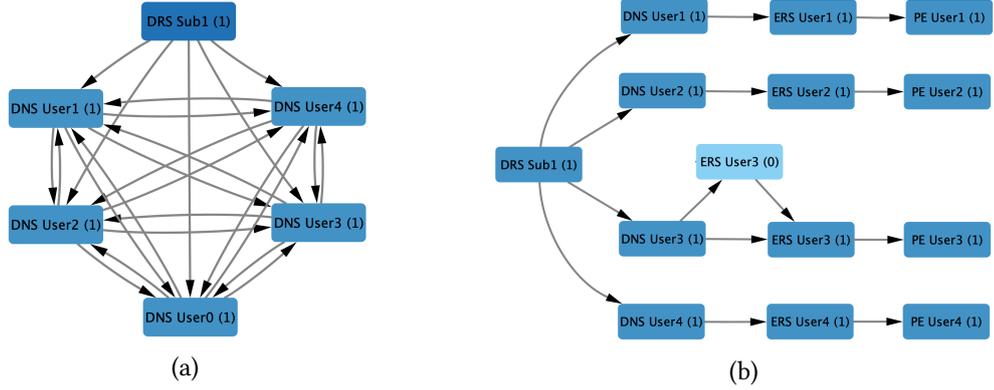


Figure 3: The action-outcome transition graphs of (a) the MeanderAgent and (b) the BLineAgent adversary in steps 1-4 of the CAGE II CybORG environment.

5. Explaining the Adversary Model

The behaviour of the adversaries is dependent on the network topology and the choice of defensive actions. In addition, there is stochasticity in both the choice and outcome of actions across all of these components. Explaining adversarial behaviour proved essential in developing effective defensive models. To better understand each adversary we, at each time step, record the choice of action, outcome and the resulting state transition. For consistency across multiple episodes we resolve IP ranges and addresses to subnets and hostnames, respectively. We observe that the connectivity (i.e., the edges) of the resulting graph provides a clear signal for differentiating the two adversaries. Figure 3 shows a subset of the observations, recorded during the first four steps of adversarial behaviour, in which the BLineAgent and MeanderAgent can be seen adopting a depth-first and breadth-first approach to attacking the network, respectively. In Section 6 we present two methods which make use of this observation to more accurately determine the class of adversarial threat than in prior work [9]. In Appendix B we include the fully extracted adversary specifications generated by our methodology.

6. Hierarchical RL Architecture

In order to improve the performance of our defensive capability we explore the use of alternative controller models. We introduce two new types of controller for this task, one heuristic and another bandit-based.

6.1. Bandit Controller Model

We employ a bandit controller that is based on the multi-armed bandit architecture. The task is to determine which of the adversaries is currently attacking the network, based on the sequence of observations. However, using a bandit or bandit-like approach comes with several challenges in this setting.

In the traditional multi-armed bandit there is no notion of state: an agent takes actions and then observes the reward. However, in the CybORG environment a unique observation cannot be used to determine the current adversary. Thus sequences of observations need to be observed and, due to the stochasticity, there are multiple sequences that can be observed over a given number of timesteps. A single bandit predicting the adversary will do no better than 50%.

This is analogous to the traditional multi-armed bandit setting. Consider the task of determining which of two slot machines has the higher payout in a casino (A): the task is trivial after several attempts. Now consider a second identical casino (B) where the payout of the machines is flipped. Again, we can find the better machine in B after some error. Finally, consider being randomly placed in A or B and having only one attempt to select the slot machine with the highest payout. As we do not know which casino we are in (as everything is identical), the best possible guess rate is 50%.

We are able to solve this problem by abstracting the observations (which casino you are in) from the bandit. In this way we define N_b bandits, one for each of the observations. As such the observation is unique to the bandit predicting the adversary. While this could also be solved by a logistic regression model, the Bandit Controller is able to learn with fewer samples, also being able to determine new adversary behaviours and learn to predict them in an online fashion.

6.1.1. Bandit Controller Implementation

The bandit learning algorithm, shown in Algorithm 1, allows the bandit controller to track the states that it has previously seen, creating a new bandit for each newly seen state. Each of these bandits is initialised with Q values for each of the actions $a \in \{0, 1, 2\}$, where these correspond to the `MeanderAgent`, `BLineAgent`, and no adversary. The Q values are updated using reward R and the number of times that prediction has been selected, $N(A)$. We train the bandit controller for 15,000 timesteps, using $\epsilon = 0.01$.

The Bandit Controller has a state different to that of its subagents. Its state is a sliding window of the last four timesteps from the CybORG environment. As we can see from Figure 3b, the minimum number of actions before an adversary has user privilege (and the first unambiguous instance of malicious behaviour) is three. A defensive agent can observe this on the fourth timestep, hence a prediction from the bandit controller only needs to happen once per episode. Finally, we use a simple reward function of +1 for a correct prediction, and -1 for an incorrect prediction.

6.2. Heuristic Controller Model

We also construct a heuristic for predicting the adversary. This approach is possible as we are able to observe the patterns that the adversaries display in a controlled version of the CybORG environment. As we can see in Figures 3b and 3a, the `BLineAgent` and `MeanderAgent` have fundamentally different strategies in the first four moves they make. Using this privileged view of the adversarial behaviour allows for a manual and formal definition of the behaviour, as defined in Heuristic 1. As in the Bandit Controller we use this heuristic once per episode, on the fourth timestep, to determine which adversary is attacking the network.

Algorithm 1 Bandit Controller Learning Algorithm.

Initialise the known states, s_n

Initialise set of bandits, B

Initialise for $a = 1$ to k :

$bandit_0.Q(a) \leftarrow 0$ // **Initialise Q values and action counter for the first bandit**
 $bandit_0.N(a) \leftarrow 0$

Predict(s):

if $s \notin s_n$:
 $s_n \leftarrow s$
 Initialise $bandit_s$
 $B \leftarrow bandit_s$
 $A \leftarrow \begin{cases} \mathit{argmax}_a(bandit_s.Q(a)) & \text{with probability } 1 - \epsilon \\ \text{random action} & \text{with probability } \epsilon \end{cases}$
 $R \leftarrow \text{prediction_result}(A)$
 $bandit_s.N(A) \leftarrow bandit_s.N(A) + 1$
 $bandit_s.Q(A) \leftarrow bandit_s.Q(A) + \frac{1}{N(A)} [R - bandit_s.Q(A)]$

Heuristic 1. *The scanning of two different hosts on the network within the first four timesteps indicates the presence of the MeanderAgent adversary. Otherwise, this is either the BLineAgent adversary or the User agent.*

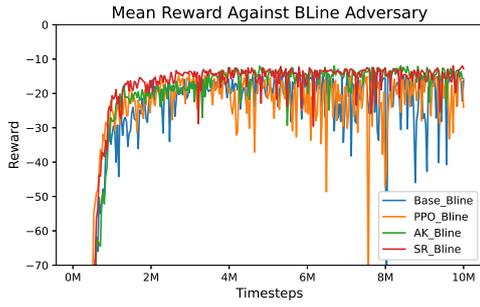
7. Evaluation

In this section we evaluate the performance of our specialist subagents against the two adversaries. We further investigate the performance of the controller models. Finally, we evaluate the full defensive model capable of defending against either adversary. We use the model described in prior work [9] as a baseline performance measure (baseline for brevity), as this has been established as state-of-the-art and achieved the best score in CAGE I. Because the scoring function assigns only penalty points (i.e., 0 is the theoretically maximum score), all the reported rewards are negative.

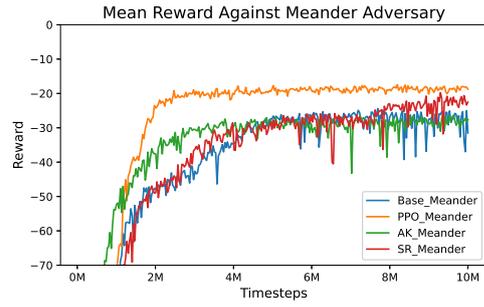
7.1. Specialised Sub Agents

7.1.1. Training Results

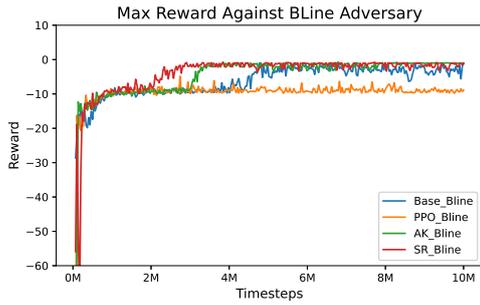
Figure 4 shows the average reward of each defensive subagent as trained against the BLineAgent (left column), and the MeanderAgent (right column). The methods of AK and SR achieve peak rewards against the BLineAgent of -12.227 and -11.465 respectively, both of which are an improvement over the baseline [9] PPO with curiosity based model, which achieves -13.475. Furthermore, removing curiosity negatively impacts the reward against the BLineAgent, as shown clearly in the max reward plot of Figure 4(c)



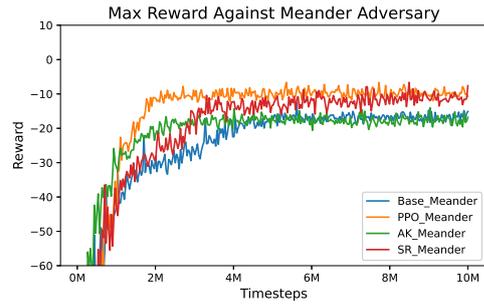
(a)



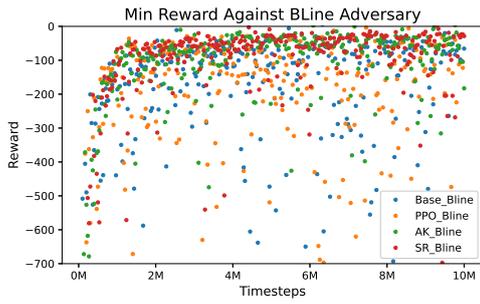
(b)



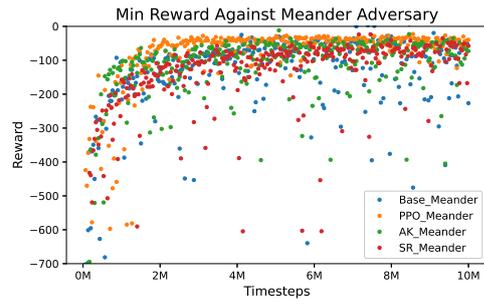
(c)



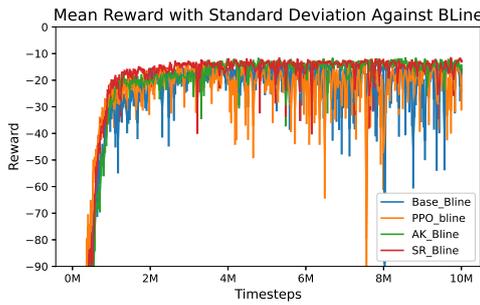
(d)



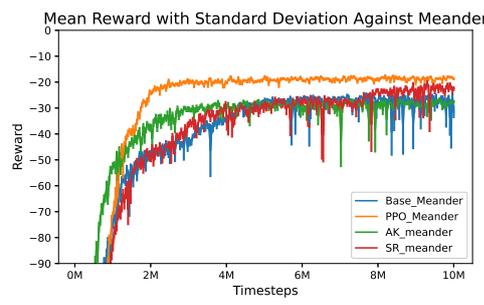
(e)



(f)



(g)



(h)

Figure 4: Mean, maximum and minimum reward of blue subagents against the BLineAgent (left) and the RedMeanderAgent (right) over 10 million timesteps.

Training Adversary	Defensive Model	BLineAgent		MeanderAgent		Mean	
		Mean Reward	Standard Deviation	Mean Reward	Standard Deviation	Mean Reward	Standard Deviation
MeanderAgent	PPO	-24.91	9.21	-17.71	5.06	-21.31	7.43
	baseline	-123.91	229.59	-30.39	47.74	-77.15	165.82
BLineAgent	AK	-14.43	30.28	-145.26	123.92	-79.84	90.20
	SR	-12.95	6.19	-269.64	235.24	-139.29	166.40
	baseline	-16.80	21.12	-201.13	143.40	-108.96	102.49

Table 1

The performance of the defensive subagents against their corresponding adversaries. Evaluated on 1,000 episodes of 100 timesteps each.

The difference in mean reward is explained by the maximum and minimum rewards. All models apart from PPO experience a first plateau in maximum reward of -9 and then step up to a second plateau of around -1. The SR agent finds the optimal policy earlier than the AK agent during training. In addition, the minimum rewards of the baseline and PPO model have greater variance than the SR and AK agents, and the AK agent has a marginally higher probability to score very poorly (i.e., below -300). Earlier optimal policy convergence and smaller policy variability makes the SR agent the best model against BLineAgent agent. This corroborates the standard deviation graph, and 1,000-episode evaluation results; in Table 1, where the SR agent displays less negative reward and with a standard deviation that is only a fifth of the AK agent.

Against the meander attacker the PPO and SR agents outperform the baseline (-24.384) with best mean rewards of -17.065 and -19.959 during training. Figure 4 shows the advantage of using a PPO 3-layer architecture which results in higher min and max rewards with reduced variance.

7.1.2. Specialist Agents

Here we evaluate the performance of our defensive subagents against their separate adversaries. We select the best performing agents from training for evaluation evaluate: PPO defence for the RedMeander and both the AK and SR defence for the BLineAgent. We evaluate each for 1,000 episodes of 100 steps and summarise our results in Table 1. For completeness, we also cross-evaluate our agents against the adversary not seen during training.

Against the RedMeander adversary, PPO defence outperforms the baseline against both adversaries resulting in a mean score of -21.3 (improvement by a factor of 3.6) and a reduction in standard deviation by a factor of more than 9. This highlights the advantage of the increased depth of the neural network over the baseline.

Against the BLineAgent adversary, we see that the SR agent is able to achieve a 1.5 times greater reward, with 4.89 times lower standard deviation. However, this comes at the cost of generality. A trend in all of the subagents is that when defending against previously unseen adversaries, the performance is significantly diminished.

Controller Agent	Prediction Accuracy	
	BLineAgent	RedMeander
PPO with curiosity (4 steps)	76.8%	0.0%
PPO with curiosity (100 steps)	30.3%	42.9%
Heuristic	100.0%	100.0%
Bandit	100.0%	100.0%

Table 2

Controller performance taken over 1,000 episodes of 4 steps, except in the case of PPO with curiosity (100 steps) which predicts the adversary at each timestep.

7.2. Controller Models

As seen in Section 7.1.2, the defensive subagents do not generalise well beyond the adversaries that they are trained against.

To address this, Sections 6.2 and 6.1 introduce two new controller architectures: Heuristic and Bandit. Here we evaluate the ability to correctly predict the adversary within the first four timesteps of an episode (as our controllers predict the adversary on the fourth timestep). For each episode, we randomly sampled one of the two red adversaries (i.e., 50% probability of selecting BLineAgent).

Table 2 shows that the baseline model has strong biases on selecting the BLineAgent agent. To investigate further, we let the baseline agent make predictions on each timestep until the end of the episode (c.f. only guessing after the 4th timestep). As seen in Table 2, the repeated guesses significantly reduced bias but accuracy remained low.

In contrast, neither our bandit or heuristic controller exhibit this bias and can perfectly predict the correct attacker type.

7.3. Hierarchical Defensive Model

Here we evaluate the complete defensive model. Table 3 reports the mean and standard deviation for the ‘best pair’ combinations of subagents as determined by our evaluation in Section 7.1 (i.e., PPO for MeanderAgent, and AK or SR for BLineAgent).

We observe that the subagents play a significant role in the improvement over the baseline. Over episodes of 100 timesteps, we are able to improve the result by at least 30% for the BLineAgent and 170% for MeanderAgent.

The lowest reward values are split evenly between the Heuristic and Bandit controllers. These models outperform the PPO controller models regardless of the subagents in four of the six combinations of adversary and episode length.

MeanderAgent performance is improved by 11.7%, which is more significant than BLineAgent (only improved by 1%) when using Bandit or Heuristic controller. Table 1 indicates that models trained with BLineAgent perform poorly on MeanderAgent. This can be explained by the fact that BLineAgent has more information about the network, so its behaviour is more predictable. In contrast, MeanderAgent’s actions have more randomness.

Controller	Subagents	30 steps		50 steps		100 steps	
		BLineAgent	MeanderAgent	BLineAgent	MeanderAgent	BLineAgent	MeanderAgent
Bandit	PPO + AK	-3.56±2.03	-6.80±1.40	-6.79±13.00	-10.10±2.30	-13.54±15.95	-17.30±4.27
	PPO + SR	-3.62±2.04	-6.88±1.42	-6.26±3.18	-10.06±2.15	-13.00±6.28	-17.56±4.51
Heuristic	PPO + AK	-3.56±2.04	-6.80±1.40	-6.79±13.00	-9.96±2.33	-14.07±27.73	-17.57±4.82
	PPO + SR	-3.71±2.09	-6.86±1.48	-6.17±3.40	-10.04±2.32	-13.06±6.14	-17.32±4.35
Baseline (PPO Controller)	PPO + AK	-4.35±2.42	-7.19±1.69	-7.45±4.27	-10.84±2.62	-14.97±8.09	-19.33±5.38
	PPO + SR	-3.95±2.18	-7.36±1.74	-6.38±3.20	-11.33±3.00	-13.14±6.45	-21.21±6.10
Baseline (PPO Controller)	Baseline (PPO subagents)	4.82±4.22	-8.78±3.21	-9.20±16.01	-19.00±20.86	-18.49±34.40	-47.60±88.16

Table 3

Performance of all subagents-controller combinations, evaluated over 1,000 episodes with a length of 30, 50 and 100 steps each.

8. Explaining the Defensive Models

It is critically important that human operators can understand the decisions made by autonomous agents. Using post-hoc XRL techniques, we determine whether our defensive agents are truly defending the network as their primary objective or as a side effect of an unintended objective. This is common in RL where agents may manipulate improperly specified reward mechanics to maximise their score in unintended ways.

8.1. Ablation Study

To understand which of the features in the observation space influence the agents decision making we perform an ablation study over knowledge of: 1) the success or failure of the previous action (hence referred to as *previous action*), 2) the adversary’s access onto a host (hence referred to as *adversary access*), and 3) whether an adversary has scanned a host (hence referred to as *adversary scan*).

The ablation results in Figure 5 show the AK and SR agents against the BLineAgent in 5a and 5b, and the PPO agent against the MeanderAgent³ in 5c. Figure 5a indicates that the AK agent’s performance is greatly affected by ‘adversary access’. While comparatively little impact seems to derive from the ablation of ‘adversary scan’ and ‘previous action’ there is some variance and the rewards fall to -812 and -539, respectively. Interestingly, the SR defensive agent is greatly affected by the ablation of the ‘adversary access’ and ‘adversary scan’, with the distribution of rewards being more negative in both cases. This is especially apparent in the case of ‘adversary scan’. Previous action has less of an effect in both AK and SR, yet still reduces the mean reward to -30.42 (a factor of 2) and -40.6 (a factor of 3), respectively. However, AK has some outlier scores that result in a minimum reward of -987.8. For PPO against MeanderAgent, Figure 5c shows that ablation of ‘adversary access’ causes a drastic reduction in reward, bringing the mean value to -781.23. Ablation of ‘adversary scan’ reduces the mean reward to -44.70, a factor of 2.52 more negative than when the observation is included.

³This defensive agent doesn’t use the previous action however we include it for completeness.

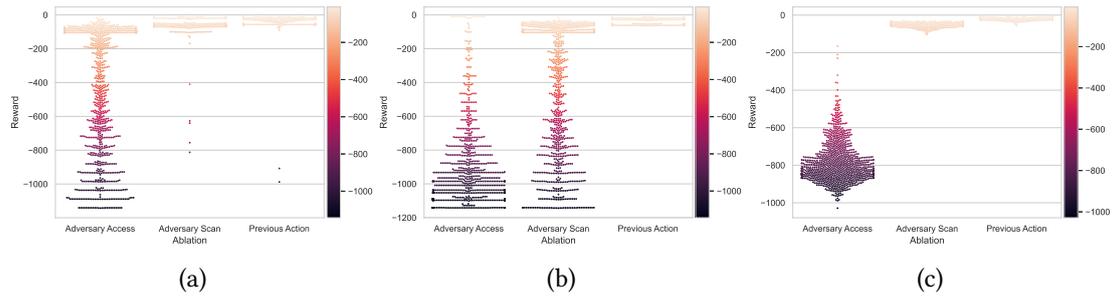


Figure 5: Ablation results for the three feature types on the reward of the (a) AK BLineAgent defence, (b) SR BLineAgent defence and (c) PPO MeanderAgent defence.

8.2. Feature Importance

To further validate the importance of ‘adversary access’, ‘previous action’, and ‘adversary scan’, we utilise a well known framework from explainable AI called SHapley Additive exPlanations (SHAP). This uses an implementation agnostic game theoretic approach to explain the importance of features in determining outputs. SHAP is able to connect optimal credit allocations with local explanations to determine SHAPley values. These values provide a way of accurately distributing the contribution of the individual features within the complete feature space [22].

Figure 6 shows the SHAPley values for the trained AK and SR subagents against the BLineAgent in 6a and 6b, and PPO against the MeanderAgent in 6c. Each point on these plots is a feature in a specific observation, with the colour representing the value of that feature.

All defensive agents observe the same trend in feature importance regardless of their training adversary: ‘adversary access’ is the most important followed by ‘adversary scan’, a trend that is also observed in Figure 5. Note that the PPO RedMeander defensive model doesn’t use ‘previous action’ and hence is not included in Figure 6c.

We show that ‘adversary access’ is an important part of the observation. This indicates that the defensive agents are aware that they need to remove the attackers from hosts. The importance is also seen in Figure 5 as the most significant shifts in reward distribution occur when ablating ‘adversary access’.

In addition ‘adversary scan’ is of importance to the agents which is clear in 5b as the defensive agent’s performance is significantly impacted in the absence of this information. This correlates with Figure 6b as ‘adversary scan’ has the greatest distribution of any of the SHAP values for the BLineAgent defensive agents. While knowledge of the ‘previous action’ has the lowest feature importance for the agents, we argue that this is still important for these defensive agents, which, with this knowledge, outperform the baseline and PPO-only models in 4. For example, take the case where a defensive agent acts to remove an adversary from a host, if this action fails then the defensive agent will have to adjust its strategy. The importance of this feature can further be seen in Figures 5a and 5b, as ablation of this feature has a non-trivial impact on the performance of the agents.

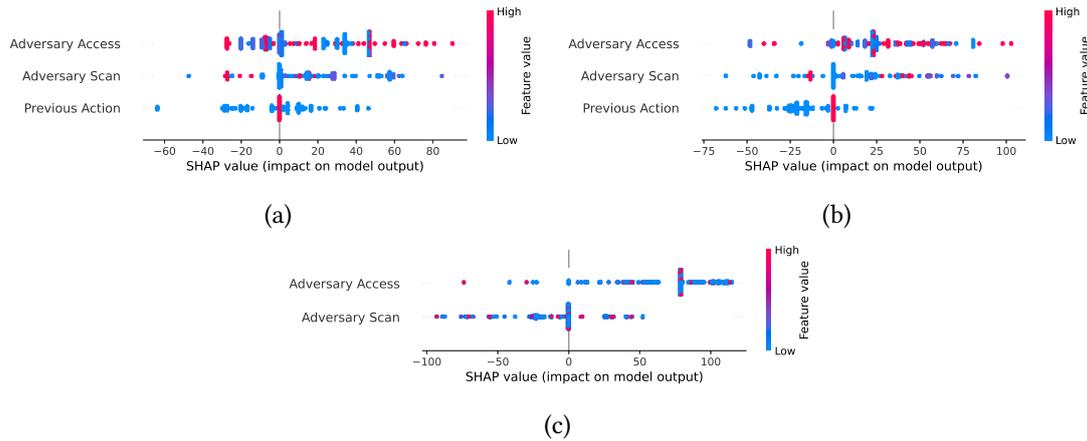


Figure 6: The SHAPley values for the three feature types on the reward of the (a) AK BLineAgent defence, (b) SR BLineAgent defence and (c) PPO MeanderAgent defence, from most to least important features.

9. Related Work

The effectiveness of RL across a range of simulated and abstracted autonomous network defence scenarios is well established in the literature. Han et al. [23] show the feasibility and resilience of RL agents under causative attacks in software defined networks. Elderman et al. [24] model network defence using the framework of a Markov game with incomplete information, highlighting the capabilities of even traditional RL methods (i.e., not DRL) in interactions between network attacker and defender. The hierarchical approach we build upon was first proposed by Foley et al. [9]. Comparatively, we propose two improved controller models based on a deeper understanding of the adversary models. We also develop improved subagents providing an explainability analysis to understand what causes the agents to defend networks effectively. Other approaches to autonomous network defence include dynamic causal Bayesian optimisation [25] as shown by Andrew et al. [26].

Several alternative network defence simulation environments have been proposed in the literature. Molina-Markham et al. [27] propose FARLAND which similarly to CybORG provides a hybrid simulation and emulation based environment capable, owing to a rich feature space, develops agents that can defend real-world networks. Microsoft have an experimental research platform CyberBattleSim [28] that offers, at a high-level of abstraction, a simulation-only network defence environment based on post-breach lateral adversary movement and system exploitation. In contrast to CybORG, CyberBattleSim places greater emphasis on credential access and data collection such as simulating a GitHub project leaking credentials in the commit history. Another simulation-only environment developed by Andrew et al. [26] is Yawning Titan (YT). Of all the network defence environments, YT offers the greatest abstraction and omits the majority of individual host details (e.g., operating system processes, network ports) needed for emulation.

RL has also been applied to several closely related problems. In penetration testing (i.e., exploitation which is a subset of the CybORG environment), Yang and Liu [29] formulate automated

penetration testing in the multi-objective RL framework and demonstrate superior performance. Independently, Tran et al. [30] explore hierarchical RL architectures for the same task based on their findings that decomposing large action spaces into smaller sets produces greater performing agents. In intrusion prevention, Hammar and Stadler [31] demonstrate that RL is capable of intrusion prevention when formulated as a multiple stopping problem. Feng and Xu [32] train a defender to protect a single device from an unknown attacker and finally, Tahsini et al. [33] use a single defender model to protect a water tank system from adversarial attacks.

10. Conclusion

Taking advantage of the rapidly increasing capabilities of neural networks and the advancements in RL algorithms, we present an improved approach to autonomous network defence. Beyond high performance, we place emphasis on the steps before and after training the model. Before training, we use a methodology to observe the adversary behaviour and inform choices in our hierarchical model. Specifically, we introduce two controller architectures, one heuristic and another bandit-based, that improve accuracy when predicting adversaries. Additionally we develop enhanced subagent architectures optimised for the specific classes of adversary. After training, our post-hoc analysis includes a feature importance and ablation study for each specialised subagent within the complete hierarchical model. Our results shed light on each agent's decision making process and help to better understand the system as a whole. This work contributes to a less studied but equally important research direction for future works in autonomous network defence.

Acknowledgements

The authors would like to acknowledge that research was partially funded by EPSRC grant EP/T51780X/1.

References

- [1] P. Speicher, M. Steinmetz, J. Hoffmann, M. Backes, R. Kunnemann, Towards automated network mitigation analysis, in: Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC '19, 2019.
- [2] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, K. Kavukcuoglu, Asynchronous methods for deep reinforcement learning, in: International conference on machine learning, PMLR, 2016, pp. 1928–1937.
- [3] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, D. Hassabis, Human-level control through deep reinforcement learning, *Nature* (2015).
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing Atari with Deep Reinforcement Learning, arXiv:1312.5602 [cs] (2013).

- [5] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal Policy Optimization Algorithms, in: arXiv:1707.06347 [cs], 2017.
- [6] O. et al., Dota 2 with Large Scale Deep Reinforcement Learning, 2019.
- [7] A. E. Sallab, M. Abdou, E. Perot, S. Yogamani, Deep Reinforcement Learning framework for Autonomous Driving, *Electronic Imaging* 29 (2017) 70–76. URL: <http://arxiv.org/abs/1704.02532>. doi:10.2352/ISSN.2470-1173.2017.19.AVM-023, arXiv:1704.02532 [cs, stat].
- [8] J. Kober, J. A. Bagnell, J. Peters, Reinforcement learning in robotics: A survey, *International Journal of Robotics Research* 32 (2013) 1238–1274. URL: <https://doi.org/10.1177/0278364913495721>. doi:10.1177/0278364913495721.
- [9] M. Foley, C. Hicks, K. Highnam, V. Mavroudis, Autonomous Network Defence using Reinforcement Learning, in: *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security, ASIA CCS '22*, Association for Computing Machinery, New York, NY, USA, 2022, pp. 1252–1254. doi:10.1145/3488932.3527286.
- [10] J. R. Williford, B. B. May, J. Byrne, Explainable Face Recognition, in: A. Vedaldi, H. Bischof, T. Brox, J.-M. Frahm (Eds.), *Computer Vision - ECCV 2020, Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2020, pp. 248–263.
- [11] C. Yu, A. Velu, E. Vinitsky, Y. Wang, A. Bayen, Y. Wu, The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games, 2022. URL: <http://arxiv.org/abs/2103.01955>. doi:10.48550/arXiv.2103.01955, arXiv:2103.01955 [cs].
- [12] T. T. Nguyen, V. J. Reddi, Deep Reinforcement Learning for Cyber Security, 2021.
- [13] X. Wu, W. Guo, H. Wei, X. Xing, Adversarial Policy Training against Deep Reinforcement Learning (2021) 1883–1900. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/wu-xian>.
- [14] R. S. Sutton, A. G. Barto, Reinforcement Learning: An Introduction, Adaptive computation and machine learning series, 2nd ed., 2018.
- [15] D. Pathak, P. Agrawal, A. A. Efros, T. Darrell, Curiosity-driven exploration by self-supervised prediction, in: *Proceedings of the 34th International Conference on Machine Learning, ICML'17*, 2017.
- [16] A. Heuillet, F. Couthouis, N. Díaz-Rodríguez, Explainability in deep reinforcement learning, *Knowledge-Based Systems* (2021). URL: <https://www.sciencedirect.com/science/article/pii/S0950705120308145>.
- [17] E. Puiutta, E. Veith, Explainable Reinforcement Learning: A Survey, in: *4th International Cross-Domain Conference for Machine Learning and Knowledge Extraction (CD-MAKE)*, 2020. URL: <https://hal.inria.fr/hal-03414722>.
- [18] M. Standen, M. Lucas, D. B., T. J. Richer, J. Kim, D. Marriott, Cyborg: A gym for the development of autonomous cyber agents, in: *IJCAI-21 1st International Workshop on Adaptive Cyber Defense*, 2021.
- [19] M. Standen, D. Bowman, S. Hoang, T. Richer, M. Lucas, R. Van Tassel, Cyber Autonomy Gym for Experimentation Challenge 1, <https://github.com/cage-challenge/cage-challenge-1>, 2021.
- [20] M. Standen, D. Bowman, S. Hoang, T. Richer, M. Lucas, R. Van Tassel, P. Vu, M. Kiely, Cyber autonomy gym for experimentation challenge 2, <https://github.com/cage-challenge/cage-challenge-2>, 2022. Created by Maxwell Standen, David Bowman, Son Hoang, Toby

- Richer, Martin Lucas, Richard Van Tassel, Phillip Vu, Mitchell Kiely.
- [21] CAGE, Cage challenge 1, in: IJCAI-21 1st International Workshop on Adaptive Cyber Defense., 2021.
 - [22] S. M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, Curran Associates Inc., Red Hook, NY, USA, 2017, pp. 4768–4777.
 - [23] Y. Han, B. I. P. Rubinstein, T. Abraham, T. Alpcan, O. De Vel, S. Erfani, D. Hubczenko, C. Leckie, P. Montague, Reinforcement Learning for Autonomous Defence in Software-Defined Networking, arXiv:1808.05770 [cs, stat] (2018). URL: <http://arxiv.org/abs/1808.05770>, arXiv: 1808.05770.
 - [24] R. Elderman, L. J. J. Pater, A. S. Thie, M. M. Drugan, M. A. Wiering, Adversarial Reinforcement Learning in a Cyber Security Simulation, in: Proceedings of the 9th International Conference on Agents and Artificial Intelligence - Volume 1: ICAART, 2017. doi:10.5220/0006197105590566.
 - [25] V. Aglietti, N. Dhir, J. González, T. Damoulas, Dynamic Causal Bayesian Optimization, in: Advances in Neural Information Processing Systems, editor = M. Ranzato and A. Beygelzimer and Y. Dauphin and P.S. Liang and J. Wortman Vaughan, volume 34, 2021.
 - [26] A. Andrew, S. Spillard, J. Collyer, N. Dhir, Developing Optimal Causal Cyber-Defence Agents via Cyber Security Simulation, in: Workshop on Machine Learning for Cybersecurity (ML4Cyber) as part of the Proceedings of the 39th International Conference on Machine Learning, 2022.
 - [27] A. Molina-Markham, C. Minter, B. Powell, A. Ridley, Network Environment Design for Autonomous Cyberdefense (2021). URL: <https://arxiv.org/abs/2103.07583>.
 - [28] J. Bono, W. Blum, Cyberbattlesim, <https://github.com/microsoft/CyberBattleSim>, 2021.
 - [29] Y. Yang, X. Liu, Behaviour-Diverse Automatic Penetration Testing: A Curiosity-Driven Multi-Objective Deep Reinforcement Learning Approach, 2022. URL: <https://arxiv.org/abs/2202.10630>. doi:10.48550/ARXIV.2202.10630.
 - [30] K. Tran, A. Akella, M. Standen, J. Kim, D. Bowman, T. Richer, C. Lin, Deep hierarchical reinforcement agents for automated penetration testing (2021). URL: <https://arxiv.org/abs/2109.06449>.
 - [31] K. Hammar, R. Stadler, Learning Intrusion Prevention Policies through Optimal Stopping, in: 2021 17th International Conference on Network and Service Management (CNSM), 2021. doi:10.23919/CNSM52442.2021.9615542.
 - [32] M. Feng, H. Xu, Deep reinforcement learning based optimal defense for cyber-physical system in presence of unknown cyber-attack, in: 2017 IEEE Symposium Series on Computational Intelligence (SSCI), IEEE, 2017.
 - [33] A. Tahsini, N. Dunstatter, M. Guirguis, C. M. Ahmed, DeepBLOC: A Framework for Securing CPS through Deep Reinforcement Learning on Stochastic Games, in: 2020 IEEE Conference on Communications and Network Security (CNS), 2020, pp. 1–9. doi:10.1109/CNS48642.2020.9162219.

Agent	Parameter	Value
AK	gamma	0.99
	network layers	[256, 256]
	Curiosity	Yes
	Beta (Curiosity)	0.2
	Eta (Curiosity)	1
	Feature Dimension (Curiosity)	53
	Learning Rate (Curiosity)	0.001
SR	Learning Rate	0.0005
	gamma	-17.710
	network layers	[256, 256]
	Curiosity	Yes
	Beta (Curiosity)	0.2
	Eta (Curiosity)	1
	Feature Dimension (Curiosity)	53
PPO	Learning Rate (Curiosity)	0.001
	Learning Rate	0.0005
	gamma	0.99
	network layers	[256, 256, 52]
Bandits	Curiosity	No
	Learning Rate	0.0005
Bandits	epsilon	0.01

Table 4
Hyperparameters.

A. Hyperparameter Values

Optimal, lower and upper bounds of the of the hyperparameters are shown in Table 4. A uniformly sampled grid search was used to determine the optimal values.

B. Extended adversary models

Here we provide the full action-outcome transition graphs for the `BLineAgent` adversary, both with and without the presence of our defensive model. Table 5 provides the definitions of all the acronyms used.

Acronym	Definition
DRS	Discover Remote Systems
DNS	Discover Network Services
ERS	Exploit Remote Service
PE	Privilege Escalate

Table 5
Acronyms used in the action-outcome transition graphs.

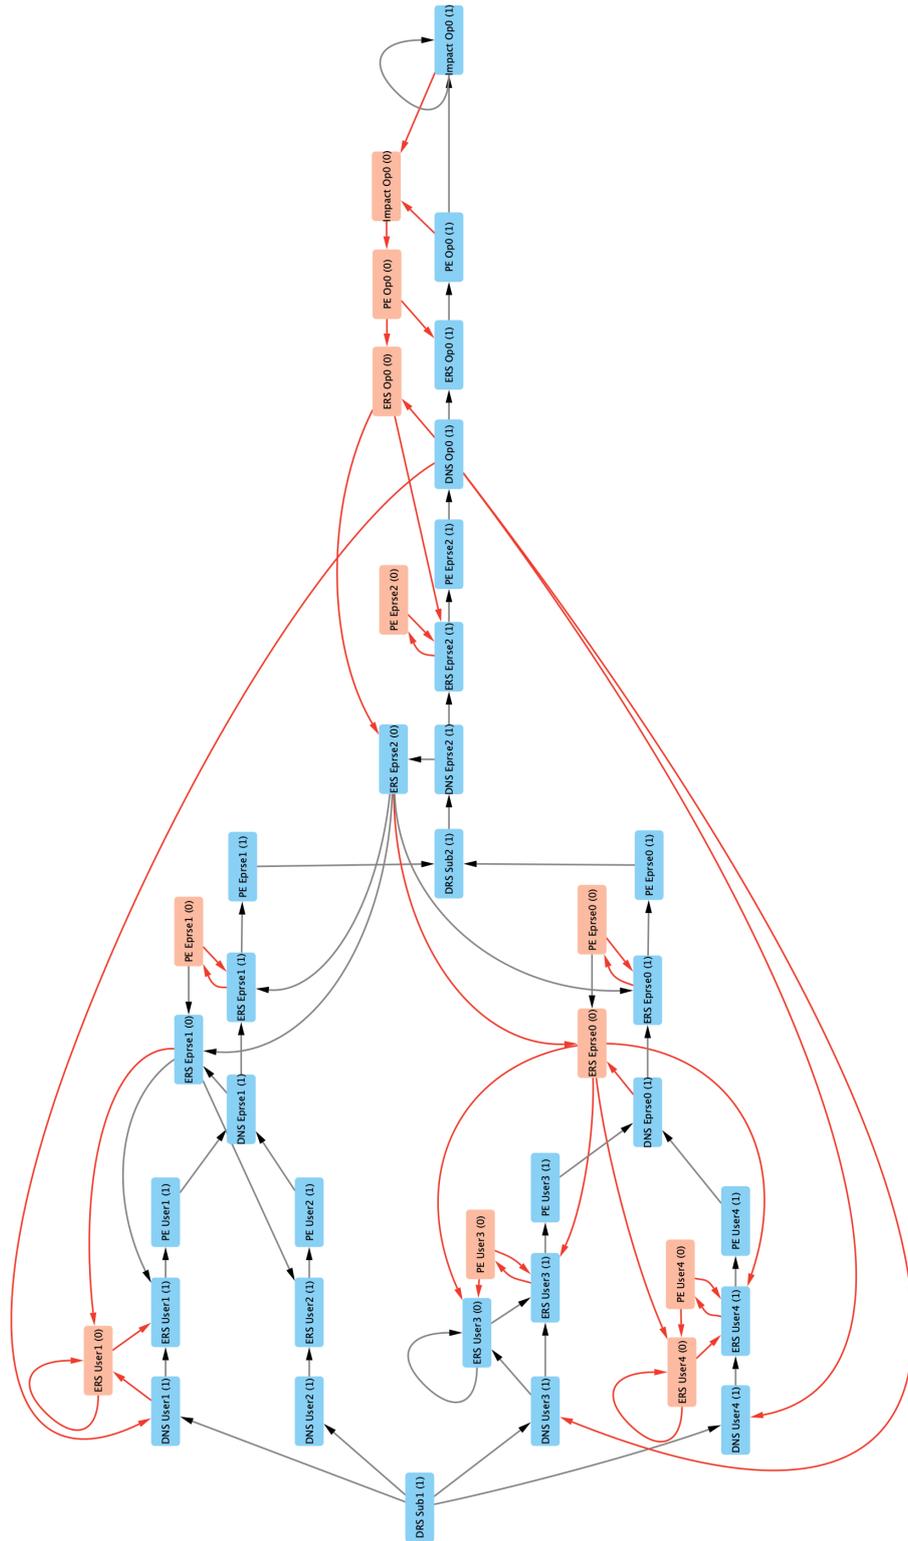


Figure 8: Action-outcome transition graph of the BLIneAgent adversary faced with our fully-trained defensive model. c.f. Figure 7, new states and transitions are shown in red.