

Applying Action Masking and Curriculum Learning Techniques to Improve Data Efficiency and Overall Performance in Operational Technology Cyber Security using Reinforcement Learning

Alec Wilson¹, William Holmes², Ryan Menzies¹ and Kez Smithson Whitehead¹

¹BMT, London, UK

²ADSP, London, UK

Abstract

In previous work, the IPMSRL environment (Integrated Platform Management System Reinforcement Learning environment) was developed with the aim of training defensive RL agents in a simulator representing a subset of an IPMS on a maritime vessel under a cyber-attack. This paper extends the use of IPMSRL to enhance realism including the additional dynamics of false positive alerts and alert delay. Applying curriculum learning, in the most difficult environment tested, resulted in an episode reward mean increasing from a baseline result of -2.791 to -0.569. Applying action masking, in the most difficult environment tested, resulted in an episode reward mean increasing from a baseline result of -2.791 to -0.743. Importantly, this level of performance was reached in less than 1 million timesteps, which was far more data efficient than vanilla PPO which reached a lower level of performance after 2.5 million timesteps. The training method which resulted in the highest level of performance observed in this paper was a combination of the application of curriculum learning and action masking, with a mean episode reward of 0.137. This paper also introduces a basic hardcoded defensive agent encoding a representation of cyber security best practice, which provides context to the episode reward mean figures reached by the RL agents. The hardcoded agent managed an episode reward mean of -1.895. This paper therefore shows that applications of curriculum learning and action masking, both independently and in tandem, present a way to overcome the complex real-world dynamics that are present in operational technology cyber security threat remediation.

Keywords

Reinforcement Learning, Cyber Security, Artificial Intelligence, Operational Technology.

1. Introduction

In previous work, the IPMSRL environment [1] was developed with the aim of training defensive RL agents in a simulator representing a subset of an IPMS on a maritime vessel under a cyber-attack. This work explores the impact of changing the difficulty of the simulator through the manipulation of values representing real world dynamics, e.g. False negative rate of alerts.

RL agents will often be significantly limited if they are not exposed to the environment in which they are intended to be deployed. This also translates to when trained agents are deployed in a real scenario. Consequently, the environment needs to replicate the real scenario as closely as possible.

This paper extends the use of IPMSRL to enhance realism including the additional dynamics of false positive alerts and alert delay. Additionally, three configurations of the environment of varying degrees of difficulty are defined and tested to understand the different levels of performance a trained RL agent can reach.

This paper also applies curriculum learning and action masking as ways to mitigate the increased levels of difficulty, showing that using these techniques is data efficient and leads to higher mean episode reward.

CAMLIS'24: Conference on Applied Machine Learning for Information Security, October 24–25, 2024, Arlington, VA

✉ Alec.Wilson@uk.bmt.org (A. Wilson); William@adsp.ai (W. Holmes); Ryan.Menzies@uk.bmt.org (R. Menzies); Kez.SmithsonWhitehead@uk.bmt.org (K. Smithson Whitehead)

ORCID 0009-0003-9181-9766 (A. Wilson); 0009-0003-9646-196X (R. Menzies); 0009-0005-2286-2140 (K. Smithson Whitehead)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Curriculum learning alone is shown to increase mean episode reward [2], [3]. Action masking alone is shown to similarly improve mean episode reward [4], [5], [6]. Action masking also has the additional benefit of significantly faster training and the ability to constrain an agent’s available action space to a set that meets user-defined criteria.

Finally, both curriculum learning and action masking are applied together. This training method resulted in the highest level of mean episode reward observed in this paper.

2. Background

2.1. Reinforcement Learning




Figure 1: Reinforcement Learning Architecture [7].

RL is a training method where an agent learns to interact with an environment to complete a task. The environment is a Markov Decision Process (MDP) and consists of a state space, action space, reward function, and a transition model [7]. At each timestep an agent will take an action on the environment. The agent will then receive a reward from the reward function and the updated state of the environment. The goal of RL is for the agent to learn to maximise the reward signal. We choose to use RL over other AI methods as it allowed the agent to learn without the need for existing datasets.

2.2. IPMSRL

IPMSRL [1] is a Gymnasium-based [8] Reinforcement Learning (RL) environment that simulates an Integrated Platform Management System (IPMS) on a vessel under a cyber-attack. An IPMS controls and monitors many ship systems across propulsion, power, steering, stability, auxiliary and ancillary systems. To achieve this, IPMS utilises a distributed control system architecture that facilitates interfaces with sensors, equipment, plants, software-based control systems and network-based data [1]. A physical representation of the IPMSRL environment has been shown in Figure 2. The configuration of IPMSRL used in this paper controls a subset of these systems and focuses on the propulsion and chilled water system. The defending RL agent receives intrusion detection alerts on components following simulated infection by a cyber attacker. These alerts are based upon the MITRE ATT&CK ICS framework¹ [9]. These alerts are then passed onto the defending agent through the observation space which the agent uses to represent the environment state, S_t , as shown on Figure 1. Subsequently, the agent chooses a discrete action, A_t contain, eradicate, recover, or wait, for a given node. Contain, eradicate, and recover represents an action space modified from NIST SP-800-61 guidance, which was adapted to an Operational Technology (OT) scenario [1], [10]. An instantaneous timestep, t , takes place and the environment produces a reward, R_{t+1} , and updated state S_{t+1} .

This sequence is then repeated until all critical nodes are compromised which would result in a negative reward or all infections are completely removed which would result in a positive reward. We also implemented an early stopping criterion of 50 timesteps.

¹© 2024 The MITRE Corporation. This work is reproduced and distributed with the permission of The MITRE Corporation.

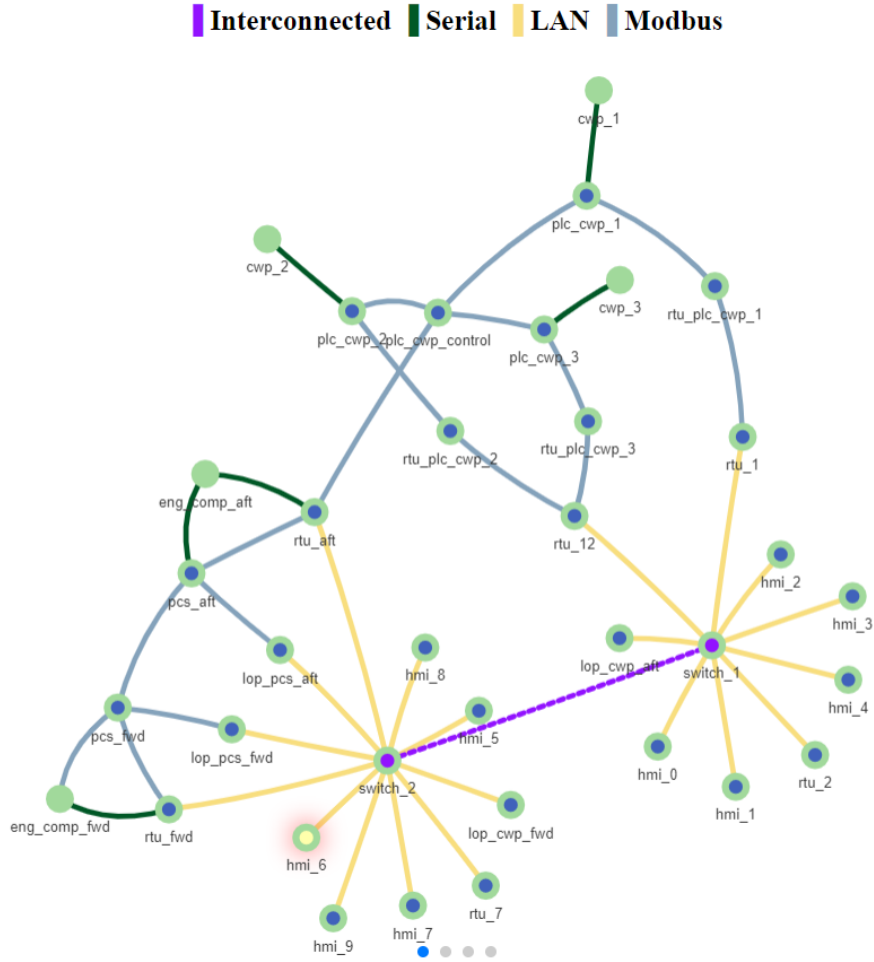


Figure 2: IPMSRL Environment [1].

2.3. PPO

Proximal Policy Optimisation (PPO) is one of the current state-of-the-art algorithms used within RL and Multi-Agent Reinforcement Learning (MARL) [11], [12]. PPO applies the policy gradient method within the actor-critic architecture. The algorithm was chosen as it is robust to hyperparameter tuning and has been shown to be performant in previous work [1], [11]. The actor, or namely the policy network, chooses the action given an observation. The critic, or namely the value network, produces an estimate of the sum of future rewards given the action from the policy network and state. This can be simplified to the actor chooses an action and the critic assesses its quality. The hyperparameters and architecture used in our experiments have been provided in the appendix.

2.4. Curriculum Learning and Action Masking

Curriculum Learning and Action Masking are two popular guided RL methods to address data efficiency concerns in RL [2], [4], [5], [13]. This paper explores both forms of guided RL for the IPMSRL cyber security environment [1], initially individually and subsequently in combination, which we show further improves performance.

2.5. Curriculum Learning

Curriculum Learning (CL) in RL is the process of increasing the difficulty of a task by periodically shaping aspects of the MDP throughout training [2], [3]. This type of guided RL is often implemented

by altering aspects of the environment such as the complexity of the transition function to increase the difficulty of optimising the reward function [4]. As a result, CL can be considered a form of transfer learning and has been shown to be beneficial for sim-to-real applications [14].

CL has been shown to reduce learning time and improve performance of the trained agent for applications including robotics and games [15], [16]. In this paper, we explore if CL offers similar benefits in the domain of cyber security by testing it on the IPMSRL environment [1]. We introduce three stages to the curriculum (Easy, Medium and Hard) where difficulty is defined as the uncertainty in the transition function. We show CL outperforms training directly on the Hard environment configuration.

2.6. Action Masking

Action Masking (AM) is a guided RL method which allows the integration of additional human knowledge into the learning process [13]. AM limits the range of actions by introducing guardrails which can prevent undesirable actions from being chosen by the agent. These guardrails require action space shaping which has similar disadvantages to reward shaping, such as increased manual set up time and susceptibility to human bias [17]. However, a key advantage of guardrails is that they can provide both increased data efficiency and user-defined constraints which are both essential considerations for cyber security applications [18].

Action masking limits the set of actions that can be executed per timestep based on the current environment state. In the context of masking for discrete spaces, this is the process of reducing the choice of actions so that undesirable or impossible actions are unavailable to the agent. Specifically, this is achieved by setting the probabilities of selecting the undesirable actions to zero or near zero during stochastic learning [4].

Previous work has shown action masking can simplify learning for the agent and result in reduced training times in video games including StarCraft II and DOTA 2 [4], [5], [6]. The focus of the work was primarily to address the high dimensionality of the action space as opposed to implementing safety critical constraints, but existing work [19] has shown how action masking can be applied in RL to improve safety for traffic-based applications.

The data efficiency benefits of action masking in cyber security applications with discrete action spaces were explored in this paper. Specifically, we mask invalid actions on the IPMSRL environment and show that the learning process provides a higher average return, as the agent focuses on learning only from the valid set of actions. This helped prevent the agent from wasting time exploring trajectories and taking undesirable actions that would not be applicable for deployment. In future work, the use of action masking could be further extended to restrict available actions based on safety-critical criteria.

We show below how action masking can also be applied to improve the realism of the IPMSRL simulation, including full details of how masking was applied in our experiments. For example, in the real world if the agent has sent a command to contain a node, then the user would likely have to wait for this process to complete before sending a different command to the given node.

2.7. Combined Curriculum and Action Masking

AM and CL often aim to achieve the same benefits of safer learning and reduced training time. We show below how these methods can be combined to improve performance over each method individually. Existing work has shown how automatic action masking can be used as a type of CL to alter the action space [20]. In contrast, we show action masking can be applied to the action space and used with vanilla CL as a method to address the uncertainty in the environment's transition function.

3. Experiments

3.1. Environment Difficulty Configurations

In a real scenario, Security Information and Event Management (SIEM) systems are used to give increased visibility of an OT system and flag any potential malicious activity. SIEMs are not perfect and suffer from False Positives (FP) and False Negatives (FN). The IPMSRL environment used in this paper has the additional feature of FP alerts.

Three defined environment difficulties were explored: easy, medium and hard: Table 1 shows the values chosen for each parameter and difficulty. The easy configuration has no FP or FN alerts, no alert delays and a 100% action success rate. The medium and hard configurations add difficulty to the environment by amending the FP, FN, action success probabilities and the alert delay. In previous work, it was demonstrated that as alert and action success probabilities decreased (FN increasing), a reduction to the agent’s performance was observed [1].

Table 1
Easy, medium and hard configurations

| Parameter Name | Parameter Description | Easy | Medium | Hard |
|--|---|---|---|---|
| Alert False Positives Probability | The probability of an alert false positive going off on each node per step | 0 | 0.01 | 0.03 |
| Alert Success Probability | The probability an alert is successfully detected on each infected node per step | 1 | 0.9 | 0.75 |
| Alert Delays (Dependent on MITRE Step) | The delay in timesteps between the node being infected/increasing infection level and the alert going off | (MITRE 1-4): 0 (MITRE 5-8): 0 (MITRE 9-12): 0 | (MITRE 1-4): 1 (MITRE 5-8): 0 (MITRE 9-12): 0 | (MITRE 1-4): 2 (MITRE 5-8): 1 (MITRE 9-12): 0 |
| Action success probability | Probability the defender action is successful | 1 | 0.9 | 0.75 |

The alerts in the IPMSRL environment are based on the MITRE ATT&CK ICS Tactics [9], with each tactic representing a different type of alert. These tactics are given a tactic level to represent them with the first tactic listed given a level of 1, with the following tactics increasing in level incrementally. The tactics are: Initial Access, Execution, Persistence, Privilege Escalation, Evasion, Discovery, Lateral Movement, Collection, Command and Control, Inhibit Response Function, Impair Process Control and Impact.

The alert delay parameter breaks the 12 tactics into 3 sections, displayed in Table ???. In the easy difficulty environment there is no delay present, in the medium difficulty environment there is a small delay to alerts earliest in the attack on a given node, the hard difficulty environment has a larger delay for the earliest stage tactic alerts and a small delay for tactics level 5-8. The rationale behind earlier stage tactics receiving a larger delay is that, generally, these tactics are likely to have a higher threshold which will need to be reached before malicious activity on a given node or network can be detected and reported as an alert, as the activity associated with these tactics is harder to differentiate from user activity. Therefore, as the environment increases in difficulty, the alert delay increases towards a more realistic scenario.

It is necessary to point out that although IPMSRL has added support for more realistic and complex dynamics, it is still an abstract representation of an IPMS and an attack on this system. The different difficulties of environment add this realism to test the performance of different training approaches and algorithms, but further work needs to be completed on IPMSRL before it can be considered representative

of a real-world system.

All of the experiments conducted in this paper use a PPO [11] based agent, trained for 2.5 million timesteps over 4 seeds with a 95% confidence interval (CI). The hyperparameters are available in the appendix.

3.2. Hardcoded Defender

A basic hardcoded defender was designed with logic developed alongside a cyber security expert. This hardcoded defender is not intended to be perfect, but in the absence using cyber security experts to act as the defender remediating the threat within the IPMSRL environment, the hardcoded defender deploys solid logic based on the NIST SP-600-61 guidance [10]. The hardcoded defender is therefore able to provide us context as to what level of mean episode reward represents a “good” performance with validated logic. The hardcoded defender takes one action per timestep, as an RL agent does in all of the experiments explored in this paper. The defensive hardcoded agent uses the following logic to determine which action to take next:

1. Contain infectable node that is connected to a critical node, with an alert.
2. Recover offline critical node.
3. Contain infectable node with an alert.
4. Eradicate infectable node that is connected to critical node, with an alert.
5. Eradicate infectable node with an alert.
6. Recover infectable node that is in a contained state.
7. Wait.

The logic follows the fundamental idea of sequentially containing, eradicating and then recovering a node, with some modifications to prioritise recovering offline critical nodes and remediating nodes that are ‘closer’ to critical nodes first.

The hardcoded defender was able to reach a mean episode reward over 10,000 episodes of 0.988 in an easy environment, 0.883 in a medium difficulty environment and -1.895 in a hard environment configuration. Figure 9 shows the comparison of the best performing RL defenders and the hardcoded defender.

3.3. Baseline Results for Vanilla PPO

The baseline results of a single agent acting in the IPMSRL environment with varying degrees of difficulty are shown in Figure 2. In the easy configuration, the agent reached near optimum performance after 1 million timesteps, winning every episode and achieving an extremely high episode reward mean value of 0.977. For the medium configuration, the agent reached an episode reward mean of 0.104 and in the hard configuration, the agent struggled to perform well, resulting in an episode reward mean of -2.791. These baseline results for vanilla PPO are all below that of the episode reward mean achieved by the hardcoded defender.

The results in Figure 3 show that there are clear challenges for the agent to learn an optimum policy using a standard training process with PPO [11] when the difficulty of the environment configuration increases. For this reason, we explored methods which enabled the data efficiency and overall performance of the agent to improve significantly.

3.4. Curriculum Learning

As discussed in the background section, CL is a training method which, in this implementation, allows the agent to initially explore simpler tasks where the agent is expected to explore more effectively, before the task is changed to a more difficult one. This enables the incremental increase in the difficulty of the environment during the agent’s training.




Figure 3: Baseline for different environment configurations.

The tasks can be changed at pre-defined points based on either standard RLib training metrics e.g. mean episode reward, or custom metrics which are calculated in the IPMSRL environment e.g. win rate. This is implemented through RLib’s TaskSettableEnv API [21].

For all the CL results reported in this paper, the number of total timesteps was used to change tasks in each training sample. The number of total timesteps was chosen to be the point at which the agent, at the current task difficulty, had reached a ‘stable’ level of performance e.g. the agent’s mean episode reward had plateaued. Therefore, for different training curricula, tasks are changed at different points depending on how quickly the agent can reach a stable level of performance. During this experiment, the tasks were changed at 850k timesteps and 1.7m timesteps. Approximately a third of the total training was completed at each task difficulty, once the task’s training had converged. The light blue lines in Figure 5 show the points at which the tasks were changed, and performance subsequently dropping sharply. Figure 4 shows the curriculum used in this experiment. Figure 5 compares the results of an



Figure 4: Curriculum for CL Experiment.

agent trained on the hard environment configuration (Figure 3), with an agent trained via a curriculum of easy, medium then hard difficulty configurations.

In Figure 5, it is clearly shown that the agent trained via CL can reach a significantly higher level of performance within the training time when compared to the baseline of “vanilla” training of a PPO agent in the hard environment configuration. The episode reward mean reached was -0.569 in comparison to the baseline of -2.791. The CL agent also significantly outperforms the hardcoded defensive agent’s performance in a hard environment, where it scored an episode reward mean of -1.895. This behaviour

is expected and supports the related literature’s conclusions [15], [16] and the intuitive premise that by leveraging the previous learning of transferable skills in simpler environments to more complex environments, an agent will perform better than attempting to learn this behaviour from scratch in a far more complex environment.




Figure 5: Baseline Hard Environment and Curriculum Learning. The changes in task are outlined by the light blue lines for the CL experiments.

3.5. Action Masking

The implementation of action masking encompasses two primary components: modifications within the environment and the development of custom models that incorporate action masking. The implementation was adapted from examples provided within RLlib [22].

Action masking utilises a binary array (mask) to identify permissible actions at each step. For this purpose, a specialised class is created, which, upon initialisation, reads a configuration file to ascertain the applicable mask conditions. This class features a method that is called at every environment step to evaluate which actions should be masked based on the current state of the environment. The observation space, structured as a dictionary, facilitates the mask’s transfer to the model, by including “action mask” and “observations” keys. The “action mask” is a 1D binary array indicating valid and invalid actions, while “observations” provide a conventional representation of the environment’s state.

For the custom model incorporating discrete action masking a PyTorch implementation is used, which is compatible with RLlib and functional for DQN and Policy-Gradient style algorithms, such as PPO [11]. This model initialises an internal fully connected network that processes solely the observation component of the observation space. In the forward pass, it computes unmasked logits by feeding the observation component through this internal network. Additionally, the action mask is transformed into an infinite mask, setting valid actions to 0 and invalid actions to a large negative value (effectively negative infinity), ensuring invalid actions are highly unlikely to be selected post-softmax application. The logits derived from the internal model are then added to this infinite mask, enabling the selection of a non-masked action. This process for action masking can be integrated with other custom models, such as a centralised critic model, allowing for both centralised critic functionality and action masking.

All instances of action masking reported in this paper used the masking conditions based on input from a cyber security expert to reflect realistic cyber defence constraints and logic. These masks are deliberately simple to avoid over-engineering and unnecessary bias. The implemented masks are:

- If there is no alert on an infectable node, mask the contain and eradicate action on that node.
- If the infectable node is contained, mask the contain action on that node.
- If the infectable node is not contained, mask the eradicate and recover action on that node.
- If an action is already in progress on an infectable node, mask all actions on that node.
- If a critical node is online, mask the recover action on that node.

Figure 6 shows the baseline results of the easy, medium and hard environment configurations when action masking is applied, compared to the baseline results without action masking. In all instances, both a dramatic improvement in data efficiency and overall performance compared to vanilla PPO can be seen. In the easy configuration, for all seeds, the agent reached an optimum level of performance in less than 100k timesteps, winning every episode, resulting in an episode reward mean of 0.977. In the medium difficulty configuration, the agent reached a good level of performance, an episode reward mean of 0.816.




Figure 6: Baseline Environments with Action Masking.

In the hard environment configuration with an action mask, the agent reached a higher level of performance than solely training on the hard environment, as shown in Figure 7, with an episode reward mean of -0.743 in comparison to the baseline hard environment episode reward mean of -2.791. The agent trained with an action mask achieved a slightly lower mean episode reward mean than CL, but was able to reach that episode reward mean at a sharper rate of learning. The AM agent, similarly, to the CL agent, significantly outperformed the hardcoded agent in the hard environment.




Figure 7: Baseline Hard Environment with and without Action Masking. The changes in task are outlined by the light blue lines for the CL experiments.

3.6. Action Masking and Curriculum Learning

Following the conclusions of the previous sections in this paper AM and CL were applied together with the subsequent curriculum shown in Figure 8.

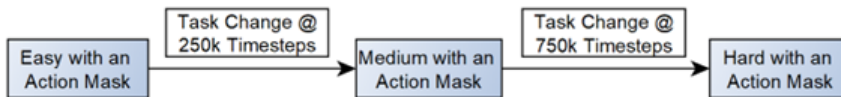


Figure 8: Curriculum for AM and CL Experiment.

The curriculum outlined in Figure 8 shows that the tasks are switched much earlier than the curriculum for the experiment without action masking presented in Figure 4. This is because, as shown in Figure 6, the learning of policies trained with action masking plateaued far sooner than policies trained without action masking. The curriculum is consequently adapted to match the attributes of agents trained with action masking, changing tasks at 250k and 750k timesteps.

Baseline Hard Environment and Curriculum Learning with and without Action Masking




Figure 9: Baseline Hard Environment and Curriculum Learning with and without Action Masking. The changes in task are outlined by the purple lines at the respective timestep totals for AM and CL experiments and by light blue lines for the CL experiments.

In Figure 9 it can be observed that when the techniques of CL and AM are combined, for all seeds, an agent trained in these conditions reached a higher mean episode reward of 0.137, in fewer timesteps, than agents trained solely with CL, AM or without either. The combination of CL and AM was therefore shown to produce the highest level of performance in comparison to the other training techniques tested and the hardcoded defender.

4. Conclusion

This paper demonstrates that the application of action masking and curriculum learning individually improved the overall performance of training a defensive agent to remediate attacks in more complex IPMSRL environment configurations. This includes real work dynamics such as false positive and negative alerts and the delay that is inherent in OT systems. The benefits of applying these techniques were even more pronounced when they were applied together.

Curriculum learning alone was shown to increase episode reward mean from -2.791 to -0.569. Action masking alone is shown to similarly improve mean episode reward. In this paper an improvement to -0.743 was seen.

Finally, both curriculum learning and action masking are applied together. This training method resulted in the highest level of mean episode reward observed in this paper, with a mean episode reward of 0.137.

As the complexity of the environment increased, the hardcoded defender struggled to maintain a strong level of performance. The hardcoded defender was able to outperform vanilla PPO, but the application of AM or CL training techniques enabled the RL defensive agent to achieve a significantly higher episode reward mean than the hardcoded defender in the hard difficulty environment.

A potential reason that the hardcoded defender struggled in the hard environment, in comparison to the best performing RL agents, is the uncertainty that a high proportion of FP alerts provides. The hardcoded agent struggles to prioritise the remediation of true alerts as there is no mechanism in its logic to establish whether an alert is a FP or not. It will have to randomly choose between which alert to remediate if there are multiple alerts active. The only prioritisation that is baked into the logic of the

hardcoded defender is to focus on nodes adjacent to critical nodes. The RL agents on the other hand may have developed policies that allow them to more efficiently decide which alerts to prioritise. This behaviour is very difficult and time consuming to encode into a hardcoded defender's logic, further displaying the benefits of using an RL based approach for autonomous cyber security.

An important note about the use of curriculum learning is the significance of defining appropriate task changing criteria. This paper implemented a simple curriculum, changing tasks at a specified number of total timesteps. There is further scope to optimise this process and potentially see higher levels of performance.

Additionally, there is a trade-off present when using action masking; the masking conditions used during training need to be present when querying the trained policy. This is a drawback in the sense that it adds a dependency to the agent's deployment, and additional bias is added through the setting of masking conditions. But the benefit of constraining certain actions which don't meet the requirements set out when developing the masking conditions is a tangible one. The use of action masking in this way therefore benefits from gains in data efficiency, overall performance, and an ability to restrict actions to meet the user-defined requirements of a given system. In future work the use of action masking could begin to consider the safety-critical nature of OT systems. Action masking could further be used to help to build trust in autonomous agents that aim to be applied to real systems.

Acknowledgments

Research funded by Frazer-Nash Consultancy Ltd. on behalf of the Defence Science and Technology Laboratory (Dstl) which is an executive agency of the UK Ministry of Defence providing world class expertise and delivering cutting-edge science and technology for the benefit of the nation and allies. The research supports the Autonomous Resilient Cyber Defence (ARCD) project within the Dstl Cyber Defence Enhancement programme.

The authors would also like to thank Lisa Gralewski, Marco Casassa Mont, David Foster, Clare Jubb, Laura Caddy, Tasha Hughes and Jake Rigby for their wider contribution to the project and paper.

References

- [1] A. Wilson, R. Menzies, N. Morarji, D. Foster, M. Casassa Mont, E. Turkbeyler, L. Gralewski, Multi-Agent Reinforcement Learning for Maritime Operational Technology Cyber Security, CAMLIS: Conference on Applied Machine Learning in Information Security 3652 (2023).
- [2] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, P. Stone, Curriculum learning for reinforcement learning domains: A framework and survey, *Journal of Machine Learning Research* 21 (2020) 1–50.
- [3] B. F. Skinner, Reinforcement today., *American Psychologist* 13 (1958) 94–99. URL: <https://doi.org/doi/10.1037/h0049039>. doi:10.1037/h0049039.
- [4] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, J. Quan, S. Gaffney, S. Petersen, K. Simonyan, T. Schaul, H. van Hasselt, D. Silver, T. Lillicrap, K. Calderone, P. Keet, A. Brunasso, D. Lawrence, A. Ekermo, J. Repp, R. Tsing, StarCraft II: A New Challenge for Reinforcement Learning, 2017. URL: <http://arxiv.org/abs/1708.04782>, arXiv:1708.04782 [cs].
- [5] OpenAI, :, C. Berner, G. Brockman, B. Chan, V. Cheung, P. Dębiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. Pachocki, M. Petrov, H. P. d. O. Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, S. Zhang, Dota 2 with Large Scale Deep Reinforcement Learning, 2019. URL: <https://arxiv.org/abs/1912.06680>. doi:10.48550/ARXIV.1912.06680.
- [6] S. Huang, S. Ontañón, A Closer Look at Invalid Action Masking in Policy Gradient Algorithms, *The International FLAIRS Conference Proceedings* 35 (2022). URL: <https://journals.flvc.org/FLAIRS/article/view/130584>. doi:10.32473/flairs.v35i.130584.

- [7] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.
- [8] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. de Cola, T. Deleu, M. Goulão, A. Kallinteris, A. KG, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. J. S. Tan, O. G. Younis, Gymnasium, ??? URL: <https://github.com/Farama-Foundation/Gymnasium>.
- [9] T. M. Corporation, ICS Matrix | MITRE ATT&CK®, 2023. URL: <https://attack.mitre.org/matrices/ics/>, publication Title: ICS Matrix MITRE ATT&CK® Type: Documentation.
- [10] N.I.S.T., NIST SP 800-61 Rev. 2 - Computer Security Incident Handling Guide, 2012. URL: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf>.
- [11] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal Policy Optimization Algorithms, 2017. URL: <http://arxiv.org/abs/1707.06347>.
- [12] C. Yu, A. Velu, E. Vinitzky, J. Gao, Y. Wang, A. Bayen, Y. Wu, The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games, 2022. URL: <http://arxiv.org/abs/2103.01955>.
- [13] J. Eßer, N. Bach, C. Jestel, O. Urbann, S. Kerner, Guided Reinforcement Learning: A Review and Evaluation for Efficient and Effective Real-World Robotics [Survey], IEEE Robotics & Automation Magazine 30 (2023) 67–85. URL: <https://ieeexplore.ieee.org/document/9926159/>. doi:10.1109/MRA.2022.3207664.
- [14] M. R. Diprasetya, A. N. Pullani, A. Schwung, Sim-to-Real Transfer for Robotics Using Model-Free Curriculum Reinforcement Learning, in: 2024 IEEE International Conference on Industrial Technology (ICIT), IEEE, 2024, pp. 1–6.
- [15] Y. Bengio, J. Louradour, R. Collobert, J. Weston, Curriculum learning, in: Proceedings of the 26th Annual International Conference on Machine Learning, ACM, Montreal Quebec Canada, 2009, pp. 41–48. URL: <https://dl.acm.org/doi/10.1145/1553374.1553380>. doi:10.1145/1553374.1553380.
- [16] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, others, Mastering the game of Go with deep neural networks and tree search, nature 529 (2016) 484–489. Publisher: Nature Publishing Group.
- [17] A. Kanervisto, C. Scheller, V. Hautamäki, Action Space Shaping in Deep Reinforcement Learning, 2020. URL: <http://arxiv.org/abs/2004.00980>, arXiv:2004.00980 [cs].
- [18] K. Thakur, M. Qiu, K. Gai, M. L. Ali, An investigation on cyber security threats and security models, in: 2015 IEEE 2nd international conference on cyber security and cloud computing, IEEE, 2015, pp. 307–311.
- [19] A. Müller, M. Sabatelli, Safe and psychologically pleasant traffic signal control with reinforcement learning using action masking, in: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), IEEE, 2022, pp. 951–958.
- [20] A. Y. Yasutomi, T. Ogata, Automatic Action Space Curriculum Learning with Dynamic Per-Step Masking, in: 2023 IEEE 19th International Conference on Automation Science and Engineering (CASE), IEEE, Auckland, New Zealand, 2023, pp. 1–7. URL: <https://ieeexplore.ieee.org/document/10260397/>. doi:10.1109/CASE56687.2023.10260397.
- [21] R. RLlib, Advanced python api's - curriculum learning, 2023. URL: <https://docs.ray.io/en/releases-2.4.0/rllib/rllib-advanced-api.html?highlight=curriculum%20leanrign#curriculum-learning>, publication Title: Advanced Python API's - Curriculum Learning Type: Documentation.
- [22] E. Liang, R. Liaw, P. Moritz, R. Nishihara, R. Fox, K. Goldberg, J. E. Gonzalez, M. I. Jordan, I. Stoica, Rllib: Abstractions for distributed reinforcement learning, 2018. URL: <https://arxiv.org/abs/1712.09381>. arXiv:1712.09381.

A. PPO Model Hyperparameters

| Hyperparameters Values | |
|------------------------|------------|
| fc net activation | Swish |
| fc net hiddens | [256, 256] |
| value net hiddens | [32] |
| lambda | 0.925 |
| kl coeff | 0.1 |
| vf clip param | 25 |
| sgd minibatch size | 100 |
| num sgd iter | 15 |
| vf loss coeff | 0.75 |
| entropy coeff | 0.0001 |
| clip param | 0.2 |
| lr | 0.0005 |
| gamma | 0.995 |
| train batch size | 5000 |
| total timesteps | 2,500,000 |
