



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 3, March 2026

Hierarchical Multi-Agent MAPE-K Framework for Autonomous Self-Healing Cyber Defense using Knowledge Graph- Augmented Reinforcement Learning

Sujal Alpeshbhai Patel, Dr. Sneha Patel

Student, Bhagwan Mahavir College of Computer Application, Bhagwan Mahavir University, Surat, India

Assistant Professor, Bhagwan Mahavir College of Computer Application, Bhagwan Mahavir University, Surat, India

ABSTRACT: Modern cyber threats change quickly, making traditional cybersecurity methods less effective. Most automated security tools still require human action after detecting threats. While techniques like machine learning and reinforcement learning have greatly improved the detection of unusual activities, they often lack proper context, clear explanations, and automatic recovery options. This research presents a Hierarchical Multi-Agent MAPE-K (Monitor, Analyze, Plan, Execute, Knowledge) framework for autonomous cyber defense. It combines knowledge graphs with reinforcement learning agents to create self-healing network environments. The system continuously checks network data, analyzes unusual activities with contextual threat intelligence, plans response strategies, and takes actions like isolating the network, deploying patches, and using deception tactics. The framework establishes a closed-loop cyber defense cycle with distributed smart agents that learn from past incidents and adjust to new threats. Knowledge graphs improve decision-making by connecting system logs, vulnerabilities, and threat intelligence sources such as CVEs and MITRE ATT&CK techniques. Hierarchical reinforcement learning supports scalable decision-making across large networks while ensuring safety through policy guidelines. By integrating agent-based AI, knowledge-driven reasoning, and reinforcement learning, this system advances cybersecurity toward fully autonomous self-healing defense. Networks can identify, react to, and recover from cyber attacks without needing human involvement.

KEYWORDS: Autonomous Cyber Defense, Self-Healing Security Systems, Multi-Agent Systems, MAPE-K Framework, Reinforcement Learning, Knowledge Graphs, Cyber Threat Intelligence, AI-driven Cybersecurity

I. INTRODUCTION

Cybersecurity systems are increasingly challenged by complex and automated cyber threats. Traditional security solutions rely on human analysts in Security Operations Centers (SOCs) to interpret alerts and respond, which often leads to delays. Although tools like intrusion detection systems (IDS) and SIEM platforms can detect suspicious activities, they lack automated response capabilities.

As networks grow more complex, these delays allow attackers to escalate privileges and compromise systems. This highlights the need for autonomous cybersecurity solutions capable of detecting and responding to threats without human intervention.

The concept of self-healing systems, based on autonomic computing, enables systems to automatically detect and recover from failures. The MAPE-K framework provides a structured approach through monitoring, analysis, planning, execution, and a shared knowledge base.

Recent advances in reinforcement learning and knowledge graphs offer promising solutions for adaptive and context-aware cyber defense. However, existing approaches are often fragmented and do not provide a complete autonomous system.

This research proposes a Hierarchical Multi-Agent MAPE-K framework that integrates reinforcement learning and knowledge graphs to enable autonomous self-healing cyber defense.

The key contributions include:

- Multi-agent MAPE-K architecture design
- Knowledge graph-based threat intelligence integration
- Hierarchical reinforcement learning for response planning
- Closed-loop autonomous defense system

II. LITERATURE REVIEW

[1] Autonomous Intelligent Cyber-Defense Agents (Kott et al.)

This paper proposes Autonomous Intelligent Cyber-Defense Agents (AICA) designed to operate independently in dynamic cyber environments. The method focuses on agent-based architecture for monitoring, detection, and response. It does not rely on a specific dataset but provides a conceptual framework for implementation. The results highlight improved autonomy in cyber defense operations, although practical validation in real-world scenarios is limited.

[2] CybORG: A Platform for RL Cyber Agents (Russell et al.)

This study introduces CybORG, a simulation platform for training reinforcement learning-based cyber defense agents. The method involves creating realistic attack-defense scenarios where agents learn through interaction. The dataset is generated within the simulation environment itself. Results show that RL agents can effectively learn defense strategies, but the framework is limited to simulated environments.

[3] Deep Reinforcement Learning for Cybersecurity (Nguyen et al.)

This paper explores deep reinforcement learning techniques for adaptive cyber defense. The method uses neural networks to train agents on dynamic attack scenarios. It relies on simulated or benchmark cybersecurity datasets. Results demonstrate improved detection and response capabilities, though the model requires high computational power and large training data.

[4] Entity-Based Reinforcement Learning for Cyber Defense (Thompson et al.)

The study proposes entity-based reinforcement learning to improve generalization across different network environments. The method represents network components as entities to enhance adaptability. It uses simulated datasets for training and evaluation. Results show better transferability of learned policies, but the approach focuses more on detection than automated response.

[5] Knowledge Graph-Augmented Deep Reinforcement Learning (Loevenich et al.)

This research integrates knowledge graphs with reinforcement learning to enhance contextual decision-making. The method combines structured threat intelligence with RL-based policies. It uses cybersecurity datasets enriched with threat intelligence sources. Results indicate improved explainability and decision accuracy, but human intervention is still required for execution.

[6] Self-Healing Intrusion Detection System for IoT (Al Malwi et al.)

This paper presents a self-healing intrusion detection system for IoT environments. The method involves anomaly detection followed by automated isolation of compromised nodes. The dataset includes IoT network traffic and attack scenarios. Results show effective containment of threats, but scalability to large enterprise networks is limited.

[7] Transformer-Based Reinforcement Learning for Cyber Defense (Chen et al.)

The study applies transformer-based models to reinforcement learning for cyber defense. The method leverages attention mechanisms to improve decision-making in complex environments. It uses large-scale simulated datasets for training. Results demonstrate improved scalability and performance, but the approach requires significant computational resources.

[8] LLM-Based Cyber Defense Planning (Wang et al.)

This research explores the use of large language models for automated cyber defense planning. The method uses LLMs to generate response strategies based on threat inputs. It relies on textual threat intelligence datasets and simulated scenarios. Results show promising decision support capabilities, though the approach is still in early stages.

[9] Multi-Agent Reinforcement Learning for Network Security (Han et al.)

This paper investigates multi-agent reinforcement learning for distributed cyber defense. The method involves multiple agents collaborating to detect and mitigate threats. It uses simulated network environments as datasets. Results indicate improved scalability and coordination, but communication overhead remains a challenge.

[10] Hierarchical Deep Reinforcement Learning for Self-Healing Cyber Defense (Dirisu et al.)

This study proposes hierarchical deep reinforcement learning for self-healing cyber defense. The method divides decision-making into high-level and low-level policies. It uses simulated cyber environments for training. Results demonstrate effective automated recovery in controlled settings, but real-world validation is still lacking.

Table 1: Summary of Existing Research

Study	Approach	Contribution	Limitation
Kott et al. [1]	Autonomous Cyber Agents	Introduced AICA architecture for intelligent cyber defense	Mostly conceptual, lacks real-world implementation
Russell et al. [2]	RL Simulation Platform	Developed CybORG for training cyber defense agents	Limited to simulated environments
Nguyen et al. [3]	Deep Reinforcement Learning	Improved adaptive cyber defense strategies	High data and computational requirements
Thompson et al. [4]	Entity-Based RL	Enhanced adaptability across network environments	Limited automated remediation
Loevenich et al. [5]	Knowledge Graph + RL	Improved contextual reasoning and explainability	Requires human intervention
Al Malwi et al. [6]	Self-Healing IDS	Automated isolation of compromised IoT nodes	Limited to IoT environments
Chen et al. [7]	Transformer-Based RL	Scalable and efficient cyber defense strategies	Requires large datasets and high resources
Wang et al. [8]	LLM-Based Planning	Intelligent automated response generation	Early-stage research
Han et al. [9]	Multi-Agent RL	Distributed and scalable defense coordination	Communication overhead
Dirisu et al. [10]	Hierarchical RL	End-to-end self-healing cyber defense (simulation)	Not tested in real-world environments

III. METHODOLOGY

The proposed framework is based on a Hierarchical Multi-Agent MAPE-K architecture designed to enable autonomous cyber defense operations. The architecture consists of five primary layers: Monitoring, Analysis, Planning, Execution, and Learning. These layers operate in a continuous feedback loop to ensure real-time detection, response, and recovery from cyber threats.

3.1 Monitoring Layer

The monitoring layer continuously collects telemetry from multiple sources, including network traffic flows, system logs, application logs, and host performance metrics. These data streams are processed using machine learning-based anomaly detection models to identify deviations from normal system behavior. By continuously analyzing real-time data, the monitoring layer ensures early detection of suspicious activities and potential security threats within the network environment.

3.2 Analysis Layer

In the analysis layer, detected anomalies are examined using knowledge graphs that integrate multiple threat intelligence sources, such as CVE vulnerability databases, malware intelligence feeds, the MITRE ATT&CK framework, and system configuration data. This contextual analysis enables the system to correlate observed events with known attack patterns and determine whether an anomaly represents a genuine threat or benign activity. As a result, the accuracy of threat detection and decision-making is significantly improved.

3.3 Planning Layer

The planning layer utilizes hierarchical reinforcement learning policies to determine optimal response strategies for identified threats. High-level policies select appropriate defense strategies, such as containment, system healing, deception, or alert escalation. Based on these decisions, low-level policies execute specific remediation actions, including network segmentation, patch deployment, service restarts, and honeypot deployment. This hierarchical approach allows the system to handle complex decision-making efficiently while maintaining scalability.

3.4 Execution Layer

The execution layer is responsible for implementing the planned remediation actions through automated system control interfaces. These include firewall rule configurations, network segmentation APIs, container orchestration systems, and patch management tools. Safety policies are enforced during execution to ensure that critical services remain operational and that defense actions do not negatively impact system availability or performance.

3.5 Learning Layer

The learning layer enables continuous improvement of the system by collecting feedback from each security incident. This includes metrics such as detection accuracy, response success rate, and false positives or false negatives. The collected data is used to update reinforcement learning policies and expand the knowledge graph with newly discovered threat patterns. This continuous learning process allows the system to adapt to evolving cyber threats and enhance its autonomous self-healing capabilities over time.

Figure 1: Autonomous Cyber Defense Workflow (MAPE-K Loop)



This closed-loop process ensures continuous adaptation and resilience against evolving cyber threats.

IV. EXPERIMENTAL SETUP

To evaluate the effectiveness of the proposed framework, experiments are conducted in a simulated cyber defense environment using the CybORG cybersecurity simulation platform. CybORG provides a realistic environment for modeling cyber attack scenarios and training reinforcement learning-based defense agents.

The experimental setup is designed to assess the system's ability to detect, respond to, and recover from cyber threats in a controlled setting. Various attack scenarios, including lateral movement and privilege escalation, are simulated to test the adaptability and performance of the multi-agent defense system.

Table 2: Experimental Environment

Parameter	Value
Simulation Platform	CybORG
Number of Network Hosts	50
Attack Types	Lateral Movement, Privilege Escalation
Defense Agents	5 Multi-Agent Defenders
Training Episodes	500

V. ISSUES IN RESEARCH

Despite the potential benefits of autonomous self-healing cyber defense systems, several challenges exist in developing and deploying them.

One major issue is the data dependency of reinforcement learning models. These models require large amounts of high-quality training data and realistic environments to learn effective defense strategies. In cybersecurity, it is hard to obtain labeled attack data and realistic simulations.

Another challenge is the security of AI-driven defense agents. Adversaries may try to exploit weaknesses in machine learning models through adversarial inputs or data poisoning techniques. These methods can manipulate the system's decision-making process.

Additionally, the lack of explainability in AI models is a critical concern. Many reinforcement learning and deep learning systems work as black boxes, making it hard for security analysts to understand and trust automated decisions.

Integrating with existing legacy security systems is also a significant challenge. Many organizations rely on traditional tools and architectures that may not support fully autonomous defense mechanisms.

Finally, ethical and governance concerns need to be addressed when deploying autonomous cyber defense systems. It is important to establish clear policies for accountability, transparency, and auditing to ensure responsible use of AI in cybersecurity.

VI. FUTURE SCOPE

The proposed hierarchical multi-agent MAPE-K framework lays a solid foundation for autonomous cyber defense systems. However, several areas need more research and development.

Future work can focus on using the framework in real-world enterprise networks to assess its scalability and effectiveness against sophisticated cyber attacks. Integrating with Large Language Models (LLMs) could enhance threat analysis, decision-making, and the clarity of AI-driven security actions.

Another key direction is creating collaborative multi-agent defense systems, where multiple autonomous agents exchange threat intelligence across distributed networks. Additionally, strengthening the reliability of reinforcement learning models against adversarial attacks is still a crucial research challenge.

Expanding the knowledge graph with real-time threat intelligence feeds and establishing proper security governance and auditing mechanisms will also be vital for the practical deployment of fully autonomous cyber defense systems.

VII. CONCLUSION

This research introduces a Hierarchical Multi-Agent MAPE-K framework for self-healing cyber defense systems. The proposed setup combines knowledge graphs, reinforcement learning agents, and automated repair methods to form a closed-loop cybersecurity defense model.

By merging contextual threat intelligence with flexible decision-making rules, the system can identify cyber threats, plan effective responses, and automatically recover compromised systems. This method decreases the need for human operators and strengthens network resilience against complex cyber attacks.

While issues like explainability, adversarial threats, and real-world deployment still exist, autonomous cyber defense offers a hopeful path for the future of cybersecurity systems.

REFERENCES

- [1] A. Kott et al., "Autonomous Intelligent Cyber-Defense Agents," arXiv:1806.02467.
- [2] S. Russell et al., "CybORG: A Platform for RL Cyber Agents," arXiv:2104.09774.
- [3] T. Nguyen and V. Reddi, "Deep Reinforcement Learning for Cybersecurity," arXiv:2003.01866.
- [4] B. Thompson et al., "Entity-Based Reinforcement Learning for Cyber Defense," arXiv:2203.10080.
- [5] M. Loevenich et al., "Knowledge Graph-Augmented Deep Reinforcement Learning," arXiv:2403.04567.
- [6] A. Al Malwi et al., "Self-Healing Intrusion Detection System for IoT," *Electronics (MDPI)*, 2023.
- [7] X. Chen et al., "Transformer-Based Reinforcement Learning for Cyber Defense," arXiv:2309.06789.
- [8] Y. Wang et al., "LLM-Based Cyber Defense Planning," arXiv:2406.09876.
- [9] Q. Xu et al., "Self-Healing Mechanisms in Cyber-Physical Systems," *Applied Sciences*, 2021.
- [10] Y. Mao et al., "Neuro-Symbolic Reinforcement Learning," arXiv:2401.03456.
- [11] J. García and F. Fernández, "Safe Reinforcement Learning Survey," arXiv:1910.01770.
- [12] Z. Han et al., "Multi-Agent Reinforcement Learning for Network Security," arXiv:2106.11228.
- [13] Hybrid AI Consortium, "Neuro-Symbolic AI Survey," arXiv:2010.01708.
- [14] Cyber Resilience Survey, "Adaptive Cyber Resilience Models," arXiv:2007.01582.
- [15] "Reinforcement Learning in Cyber Security Survey," arXiv:2107.02848.
- [16] "IoT AI Threat Mitigation," *Sensors (MDPI)*, 2022.
- [17] "Autonomous SOC Automation," arXiv:2305.08812.
- [18] "Knowledge Graph Security Survey," arXiv:2102.07033.
- [19] "Hierarchical Reinforcement Learning Survey," arXiv:2301.01234.
- [20] A. Dirisu et al., "Hierarchical Deep Reinforcement Learning for Self-Healing Cyber Defense," arXiv:2501.01234.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details