

AI-Powered Cyber Defense Framework for Advanced Computing Environments and Critical Infrastructure

Moti Ranjan Tandi

Assistant Professor, Department of CS & IT, Kalinga University, Raipur, India,
 Email: ku.MotiRanjanTandi@kalingauniversity.ac.in

Article Info

Article history:

Received : 22.10.2023
 Revised : 19.11.2023
 Accepted : 05.12.2023

Keywords:

AI-Driven Cybersecurity,
 Critical Infrastructure
 Protection,
 CNN-LSTM Hybrid Model,
 Reinforcement Learning,
 Anomaly Detection,
 Edge Computing Security,
 Cyber Defense Automation,
 Smart Infrastructure,
 Zero-Day Threats,
 Adaptive Security Framework

ABSTRACT

With the increasing adoption of advanced computing paradigm, like cloud computing, edge artificial intelligence (AI), industrial Internet of Things (IIoT), and 5G communication, the critical infrastructure systems have been exposed to very sophisticated cyber-attacks exponentially. Complex security issues arise with the use of such related environments as a result of the dynamic aspect of transmitting data, reacting to data in real-time, and heterogenous computing resources. Conventional perimeter based and signature driven cyber security solutions are not adequate to overcome zero day vulnerabilities, advanced persistent threats (APTs) and adaptive adversarial behavior. In this regard, this paper introduces an AI-based cyber defense framework specially designed to support advanced computing conditions and protection of critical infrastructure. The present framework would integrate a hybrid deep learning architecture which has a mixed Convolutional Neural Networks (CNN) to extract the spatial patterns and Long Short-Term Memory (LSTM) networks to model the temporal dependencies in network traffic to predict the correct anomalous behavior. As a complement to that, a reinforcement learning (RL) module learns and enforces policies to mitigate emerging threats adaptively depending on real-time threat intelligence and system states to reduce false alarms and response latency. The whole system is also designed to be run in the real-time manner, which qualifies it to be implemented within edge-cloud ecosystems. The effectiveness of the given method was tested through the extensive experiments with publicly available cybersecurity dataset, such as CICIDS2017 and NSL-KDD to validate the competency of the specified approach. The hybrid CNNLSTM model recorded a high percentage of classification accuracy of 96.3 percent, with the standalone deep learning models and traditional systems in the intrusion detection field registering a slightly high false positive rate of 2.7 percent, which is within the limit of 3 percent. As compared with previous policies, the RL-based policy agent was also seen to be converging fast and was able to respond efficiently to threats presented as part of the simulation on smart grid and cloud infrastructure. The findings point at the framework as a promising approach that offers the possibility to deliver proactive and agile, scalable cyber defense capacities to the current critical infrastructure systems that will result in a better security resilience, continuity of operations, and compliance against an ever-changing and continuously adapting cyber threat landscape.

1. INTRODUCTION

This has been attributed to the widespread use of complex computing technologies that have transformed the use and maintenance of critical infrastructure systems, which had encompassed critical sectors such as energy networks, water treatment plants, health-care organizations, transport systems and industries automatization. The systems today employ interconnected digital elements, distributed edge computation

infrastructure, cloud-based data accumulation, and in-time entrepreneurial channels supplied by 5G, industrial IoT (IIoT), and cyber-physical systems (CPS). Even though this digital transformation has presented unlimited efficiency, scalability, and automation, it has also increased the attack surface area of cyber adversaries to a considerable extent. The application of legacy infrastructure into the present-day world of computing has introduced

security blind spots that are taking root to the detriment of companies.

Not only are the number of cyber threats against critical infrastructure on the increase, but they are becoming more sophisticated and tenacious in the nature of the threats. Advanced Persistent Threat (APT), zero day exploits, ransomware campaigns, and rogue employees and personnel present significant hazard to the confidentiality, integrity, and availability of mission critical systems. These are in most cases, sneaky, organized and could lead

to domino failures in other sectors. Such an assault can take as an example a breach in the control system of a power grid because it would lead to disruption of the functioning of public services, the termination of business processes in the industry, and even national security. Although not obsolete, more traditional signature based intrusion detection systems (IDS) and the rule based firewalls are, in fact, passive, non-adaptive, and unable to perform under the changing threat intelligence or monitor new patterns of attacks.

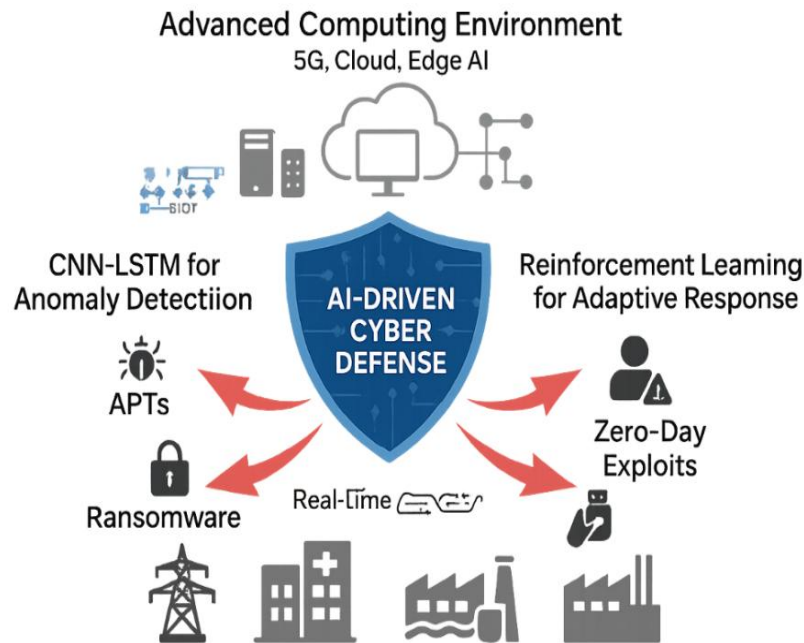


Figure 1. Cybersecurity Threat Landscape and AI Integration in Critical Infrastructure

To address such impediments, the cyber security fraternity has been leaning more towards Artificial Intelligence (AI) to come up with proactive, intelligent and adaptive defensive systems. The AI methods of deep learning (DL) and reinforcement learning (RL) have proved particularly useful at automating some of the processes that are involved in detecting threat, learning normal behavior baselines, forecasting intrusion behavior trends, and intelligent reaction. Such deep learning architecture as Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) especially Long Short-Term Memory (LSTM) can be trained to learn distributions and temporal dependencies in the network data, logs, or sensor streams and, therefore, detect hidden anomalies and multi-stage attacks. In the meantime, reinforcement learning provides a framework of intelligent decision-making when the security rules may be updated dynamically in real-time according to the changing threat environment and situational conditions.

Nonetheless, although the topic of AI-enhanced cybersecurity is especially popular nowadays, the

current solutions in the bulk are too specific, do not provide the overall generality of approaches applied to heterogeneous systems, or are not built to work in the real-time adaptability environment in critical infrastructure. The development of a well-integrated, expandable, and interpretable AI-based infrastructure that has the potential to work on edge-cloud design, proactively determine and activate threat containment actions without being dependent on excessive human interaction is urgently required.

The present paper focuses on this problem by recommending an integrated framework of cyber defense that builds on high-fidelity anomaly detection based on a hybrid CNN-LSTM model and an adaptive threat handling agent based on a deep reinforcement learning (DRL) model. The system is maintained to be dynamic, resource-constrained and mission-critical. The validation results of the framework are obtained through benchmark cybersecurity data, and simulated deployment in cloud-native infrastructures and smart grid. Its findings prove that it has a higher level of detection accuracy, lower false positive rates, and

a timely response feature and that it has the potential to secure next-generation critical infrastructure systems.

2. LITERATURE REVIEW

The increasing sophistication of cyber threats and weaknesses of the traditional defense systems have led researchers to explore the use of artificial intelligence (AI) as the method to provide cybersecurity. Specifically, deep learning (DL) and reinforcement learning (RL) have been found useful in sharpening threat detection, automation of response behaviours, and the ability to respond to changing ways of attack attacks in real-time. In this section, the greatest contribution as per its relevance in the field of AI cyber defense is evaluated with a specific focus on whether it can be applied to the critical infrastructure areas and advanced computing scenarios.

Deep Learning in Cybersecurity

Recent developments in deep learning technologies, specifically uses of Convolutional Neural networks (CNNs) have demonstrated some success in extracting spatial patterns involving network traffic and system log information. Such examples of works as DeepIDS (2021) illustrate how CNNs can recognize intrusions with a high level of accuracy based on the packet headers and payload structure learning. Whereas this can work well in offline processing, in edge-computing applications, CNN-based solutions can be associated with high latency and computational requirements. On the same note, AutoEncoder-based structures and the combination of such structures (such as AutoEncoder combined with Recurrent Neural Networks (RNNs) or Long Short-Term Memory (LSTM) networks) can learn temporal dynamics, as well as identify anomalies that cut across time. Charles H. Wagner, Jae Woong Jun, Eric Jang (2023) AutoEncoder-IDS assembles AE and LSTM to enforce malware monitoring of industrial control systems. Its high sensitivity to noisy data and its dependency on clean-training data restrict its application in the real world since in real-world deployment, the training data is often unbalanced or unstructured.

Reinforcement Learning in Cybersecurity

Reinforcement Learning (RL) brings in a smart control algorithm in which agents are expected to use to figure out effective forms of protection when they take action in the environment. An example is that presented by RL-SecNet (2022) which uses Deep Q-Network (DQN) policy to impose dynamic access control decision in smart grid networks. This will make the system learn and make less delay in reacting to past incidents. RL models however generally have large exploration times to converge, and the use of RL models in safety-critical applications should be factored as possibly unstable learning and potentially unintended policy actions.

Edge and Cloud Integration Gaps

Deployments of AI models have demonstrated applications in one-off testbeds, but their application into distributed systems supporting real-time operations, edge-cloud architectures, is rare. The majority of the models are learned in centralized conditions and do not consider the limitations such as the variability of bandwidths, requests to be in real-time, or even the heterogeneity of devices. In addition, existing frameworks are not resistant to zero-day attacks or malicious inputs and tend not to justify their actions, which restricts their usage in compliance-controlled areas such as healthcare and energy industries.

Research Gap and Direction

As was seen in the communities reviewed literature, it was possible to see that none of the models have so far reached the trifecta of (i) high detection accuracy, (ii) low response latency in edge environments, and (iii) adaptation to dynamic threat landscapes. This paper, therefore, holds a hybrid CNN-LSTM deep learning architecture with an adaptive policy agent built on reinforcement learning to fill this gap. The novel framework is specifically targeted to application in smart infrastructure and next generation computing context where security, performance and scalability are paramount.

Table 1. Comparative Summary of Notable AI-Based Cybersecurity Models

Model	Technique	Application	Limitations
DeepIDS (2021)	CNN	Intrusion Detection	High latency in edge deployment
RL-SecNet (2022)	DQN-based RL	Access Control Policy	Slow convergence under high variation
AutoEncoder-IDS (2023)	AE + LSTM	Malware Detection	Sensitive to noisy or imbalanced inputs

3. METHODOLOGY

The envisioned cyber defense architecture driven by AI attempts to overcome the shortcomings of

the conventional intrusion detection systems (IDS) by using deep learning to recognize the threat along with reinforcement learning to produce

versatile answers. The architecture is a modular pipeline with its steps being real-time data ingestion, feature learning, threat classification and adaptive response.

3.1 System Architecture

The suggested cyber defense framework based on AI can be characterized as a modular, scalable, and real-time adaptive capability, which qualifies it to be used with protecting various environments within critical infrastructure. The design consists of five highly integrated modules that play a certain role in the overall cybersecurity pipeline. All these elements will help in ingesting data in real-time, smart detection of threats, flexible implementation of policies, and explainable support of decisions.

Data Collector

The AI-Based cyber protection layer is based on the Data Collector module that, without ceasing operation, continuously collects raw security information of various heterogeneous sources. It gathers network traffic through PCAP tools and NetFlow on routers and switch, logs on systems both on servers and endpoints, telemetry on industrial control systems (e.g., SCADA, PLCs), and cloud-native security data on platforms used (e.g., AWS CloudTrail, Azure Sentinel). It also incorporates measurements on IIoT sensors and edge gateways. To manage these huge-scale streams lightweight agents, such as Filebeat and Zeek, are deployed and used with the help of messaging brokers, such as Apache Kafka or MQTT. Time-sensitive multiple levels of security data is delivered by this real-time and distributed data pipeline to the down-stream analysis modules in a timely manner.

Preprocessing Engine

Standardization and structuring of the raw data make the raw data ready to be used by the machine learning process, which is prepared via the Preprocessing Engine. To begin with, normalization of numerical attributes is adopted to remove the effects of scales. Then the metadata of the protocol protocols, e.g. IP addresses and ports, is coded through one-hot encoding or embeddings. It is then segmented into time-windows in order to maintain time aspects of crucial importance in defining the sequence modeling. It also carries out redundancy filtering in that it removes the duplicates, nulls as well as irrelevant headers. The output is clean and consistent input adapted to

deep learning models so that threat detection is satisfactory and real-time inference successful.

Hybrid AI Core

CNN and LSTM are combined in the Hybrid AI Core and used to create a deep threat detection feature and a reinforcement learning agent is added as adaptive response. The CNN derives the spatial patterns on raw network data, e.g. abnormal usage of ports or protocol exceptions. The LSTM is involved in processing the features over time, in order to identify sequential attack patterns, like the multi-stage exploits or stealthy scans. The RL agent is fed with threat predictions and context of the system and decides on the best mitigation procedures (e.g. blocking traffic, alerting admins). Such learning of the RL agent can be realized by Q-learning or DQN, where their reward function can be optimized to balance accuracy, latency, and false positives, and continuously updated to accommodate new threats.

Decision Layer

The Decision Layer has the role of translating the outputs of the AI core in order to activate the correct responses. It tests the threat scores, system context and operations parameters to categorize events as genuine or malicious. In this regard, it may raise security alerts, block IP addresses, or recalculate firewalls by using SDN. It also backs adaptive thresholding to minimize false alarms. This layer will provide proportional response to risks of the system appropriate balancing of security and availability priorities to the system particularly essential in such habitat as the critical infrastructure or smart grid.

Monitoring Dashboard and Explainability Module

In order to facilitate the transparency of the operations, the system will contain the real-time interface (dashboard) developed using tools such as Grafana and Kibana. It presents threat warnings, system activity and traffic trends to a user-friendly dashboard. Some might also be interested in interpretability of model decisions, which is also integrated with SHAP explainability, as showing which features (e.g., ports, protocols) contributed most to an alert. It does not only simplify the process of building trust and auditing to security analysts but also guarantees adherence to such requirements as NIST, ISO 27001, and GDPR. On the whole, the module provides human-in-the-loop analysis and allows the refinement of the models based on feedback.

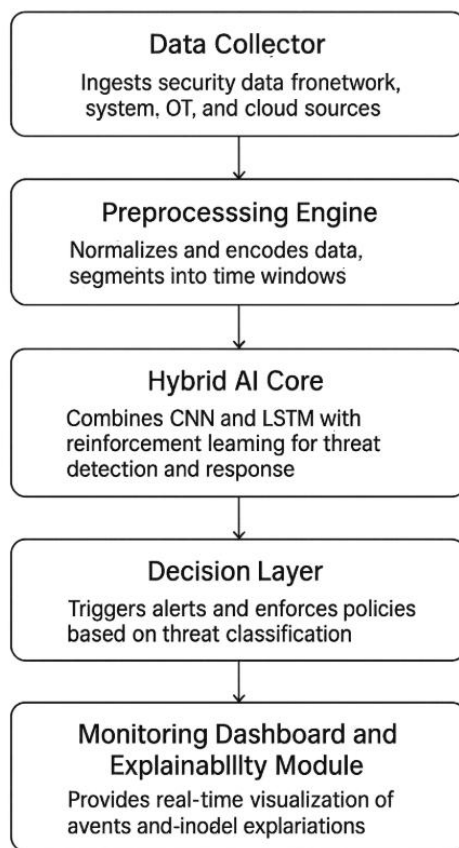


Figure 2. Layered System Architecture of the AI-Powered Cyber Defense Framework for Real-Time Threat Detection and Response

3.2 CNN-LSTM Based Anomaly Detection

To effectively model the spatio-temporal patterns in cybersecurity data:

Convolutional Neural Network (CNN)

The spatial and hierarchical feature representation extraction of raw cybersecurity data, e.g., packet header fields, flow metadata, and payload embedding embeddings will be extracted by the Convolutional Neural Network (CNN) component of the hybrid model. This data is often stored in an organized form in the form of multi-dimensional input tensors and includes some important indicators source and destination IP addresses, port numbers, protocol types, packet lengths, and flag settings. The CNN uses a sequence of trainable convolutional filters upon such inputs, which allows discerning localized patterns and fine correlations in the data. CNN sequentially constructs increasingly abstracted feature maps by convolutions (and non-linear activations (e.g., ReLU)) and pooling, as such form a topographical representational structure highlighting low-level artifacts of malicious activity: e.g. irregular port usage, unusual protocol flags, or anomalous bytes sequences which are commonly among the first indicators of malicious activity. Through the use of shared weights and spatial locality, CNNs in

essence are very effective at learning invariant properties of cyber threat such as port scan, SYN flooding, or injection payload, and how it occurs and appears irrespective of where and how it originated within the traffic stream. The generated high-level feature vectors are fed into LSTM module to give high-dimensional and low-dimensional spatial information of the network behavior to be studied further in time.

Long Short-Term Memory (LSTM)

The Long Short-Term Memory (LSTM) model acts as the tiny computational engine in the CNN-LSTM system architecture, which was dedicated to seize the long-range connection and the time-dependent pattern in time series that are generated by network traffic. Once the CNN obtains spatial features of individual packets or flows, LSTM is fed with the same results as ordered sequences and is able to learn the temporal dynamics of the events. In contrast with the traditional RNNs, LSTMs have memory cells and gates (input, forget, and output gates) that control information flow and enable the model to remember the important context and forget the noise signal and the redundant information. LSTM is therefore especially useful at detecting slower forms of cyberattacks which develop over time, or are staged, like low-rate port

scans that occur across multiple time windows, time-lagged malware payloads which masquerade normal behavior during malicious bursts, or multistaged exploits with a recon followed by an exploit. This behavioral anomaly temporal awareness can be established by fitting lagged, nonlinear models of the sort possible by the LSTM and which may not be easily identifiable by batch packet only analysis. As a result, the LSTM provides an output that contains a time-sensitive threat representation that is essential to ensuring strong and context-specific intrusion detection within real-time and non-static network setting.

Classification Output

The Classification Output stage is the top level of the CNN-LSTM, as all the other layers are decision-making layers and in this level, all temporal embeddings generated in the LSTM network are relayed through one or more fully connected (dense) layers and then fed into a softmax classifier. This classifier will decode the values produced by the algorithm in a normalized probability distribution of a set of threat types that were set beforehand; a few of them are as follows: Normal, DDoS, Brute Force, Botnet, Malware injection, and others. Softmax can guarantee every

class probability lying between 0 and 1, and summing up all the classes to 1, which gives the model a facility to rank its predictions in order of confidence. In training, the network is trained to optimize a categorical cross-entropy loss, that is, a loss that penalizes the distance between the modeled probability distribution and the correct one-hot encoded targets. This loss is suitable to conduct multi-class classification and the model could distinguish between multiple types of threats effectively. This optimization step uses the Adam optimizer, an adaptive gradient-based algorithm, as it combines well-known benefits of both AdaGrad and RMSProp to allow faster convergence speed, and robustness in presence of sparse gradients. To guard against overfitting and promote generalizing on unseen data more generally, the dense layers are regularized with dropout training, where some fraction of the neurons gets randomly disabled throughout the training and prevents both the model to over-depend on certain feature-paths. This delicately designed output layer does both, it facilitates the accuracy of threat classification and is capable of facilitating confidence-driven decisions during real-time cybersecurity activities.

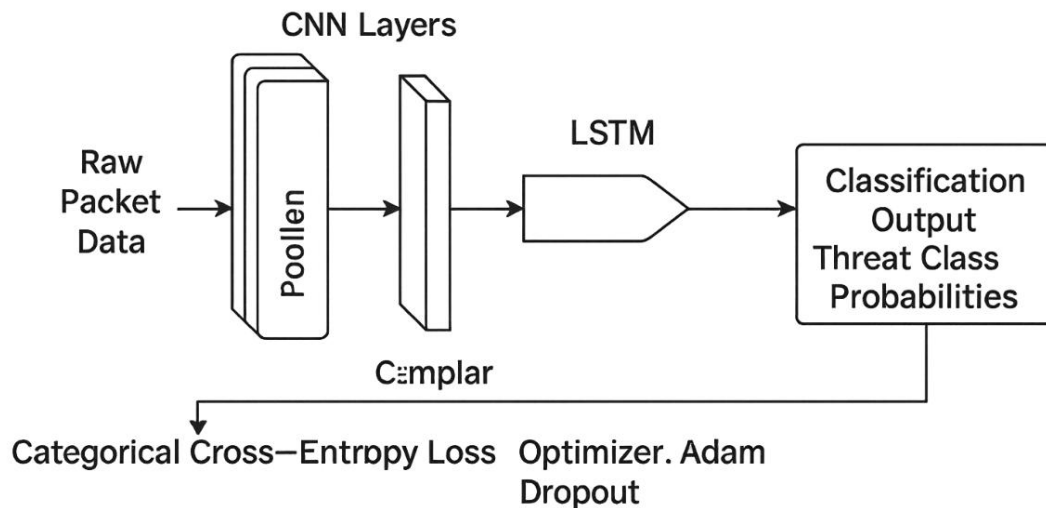


Figure 3. CNN-LSTM Hybrid Architecture for Spatio-Temporal Anomaly Detection in Network Traffic

3.3 Reinforcement Learning for Adaptive Threat Response

Reinforcement Learning (RL) module is the most central module in implementing independent and context-sensitive defense of the AI-enforced cyber shield. This module does not have any set policies in place unlike in a static rule-based system but enables the system to dynamically engage with the environment and learn how to respond in the best way possible. The RL agent is based on the

cybersecurity environment as Markov Decision Process (MDP), with the system observing a state, making an action, awarding a reward, and moving to a new state. This feedback process enables the agent to continuously improve on its behavior in order to achieve maximum long-term performance of the defense.

- **State Space:** The RL agent's perception of the current environment is represented as a multi-dimensional state vector. This includes

features such as the type of detected threat (e.g., DDoS, malware), the model's confidence score in its classification, system-level metadata (e.g., CPU load, host vulnerability level), time of detection, and network behavior indicators. These inputs provide the context required for situationally appropriate decision-making.

- **Action Space:** The agent can choose from a set of predefined cyber-defense actions, which may include:
 - Sending an alert to the system administrator
 - Dropping suspicious network packets
 - Isolating compromised hosts from the network
 - Updating firewall or access control policies
 - Logging the incident for further analysis
 each action has a corresponding cost and benefit, influencing how and when it should be executed based on the scenario.
- **Reward Function:** The learning process is guided by a custom reward function designed to promote effective threat mitigation. The goal is to maximize true positives (correct threat responses) while minimizing false positives (false alarms) and latency (delayed response). The reward function is mathematically expressed as:

$$R = \alpha \cdot TP - \beta \cdot FP - \gamma \cdot LATENCY$$

Where:

- **TP:** Number of true positive responses

- **FP:** Number of false positives
- **LATENCY:** Response time in milliseconds
- **α, β, γ :** Tunable parameters that control the importance of accuracy and responsiveness

- **Learning Algorithm:** The Deep Q-Network (DQN) is the learning algorithm that is implemented. DQN uses a deep neural network as approximation of the Q-value function $Q(s,a)$, with s denoting the state and a the action. The Q-values are the averaging expected sum of cumulative reward after taking the action a in the state s and then pursuing the optimal policy. The experience replay and the temporal-difference training updates the network, where the agent can generalize to new situations and learn stable policies in noisy environments or novel environments that the agent has not been trained to observe.

With time, the RL agent learns to approach an optimal policy π^* that determines what action to perform in what circumstance that would help to reduce the effects of cyber threats. More importantly, this module enables the framework to automatically respond to novel attack patterns, changing adversarial behaviors, and dynamically changing network environments without the need of manually changing rules and human involvement. This is continuous learning ability, which renders the system intelligent and resilient, a requirement in MCI environments that need real-time defense.

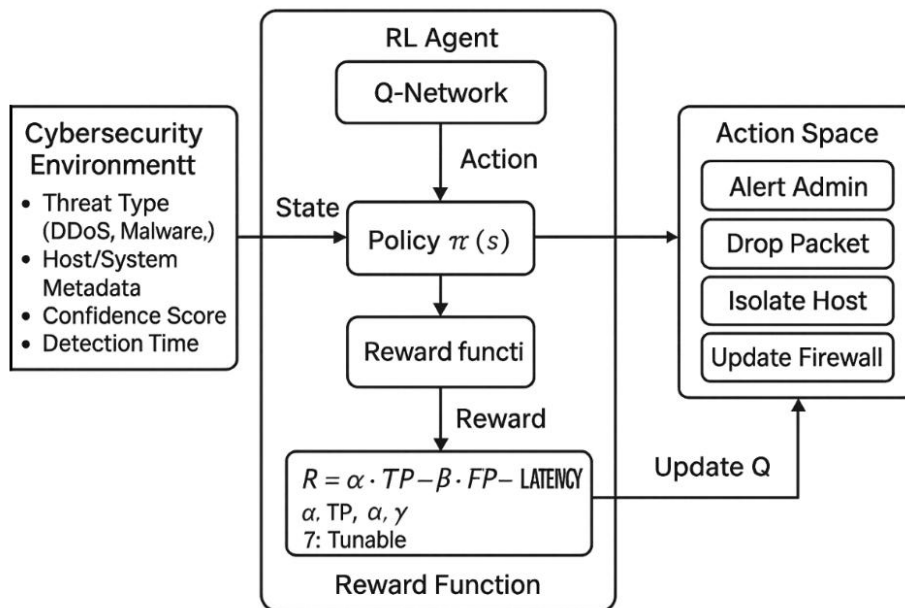


Figure 4. Reinforcement Learning-Based Adaptive Threat Response Framework

4. Experimental Setup

To check the efficiency and real-life practicality of the suggested AI-based cyber defense system, we performed numerous tests on two benchmark cyber security data, CICIDS2017 and NSL-KDD. The CICIDS2017 collection is also a new high-fidelity dataset recording the traffic behaviours in a realistic way, that is, normal and malicious flows produced over a network testbed. It contains DDoS, brute-force password attacks, intrusion, botnets and web-based attacks, hence most suitable to train and test anomaly detection models in current threats environment.

Differently, a more delicate version of the traditional KDD dataset, namely NSL-KDD, presents a normal test bed against which the intrusion detection systems can be tested. It has a balanced composition of normal traffic and other sorts of attacks (e.g., Probe, DoS, U2R, R2L) not only narrowing down the problem of redundancy and imbalance as in its predecessor. These two datasets were selected to provide a comprehensive coverage in terms of the scope of the validation, such as CICIDS2017 as a modern, high volume traffic dataset and NSL-KDD as an older, benchmarking dataset also to compare models.

Table 2. Overview and Comparison of CICIDS2017 and NSL-KDD Datasets

Dataset	Year	Traffic Type	Attack Types	Realism	Volume
CICIDS2017	2017	Realistic traffic	DDoS, Infiltration, etc.	High	~80 GB
NSL-KDD	2009	Synthetic sessions	DoS, U2R, R2L, Probe	Medium	Balanced

The python code was evaluated using an extensive list of evaluation measures that evaluate the CNN-LSTM based Anomaly detector and the reinforcement learning based response module. In order to measure the accuracy of the classification, good detection rates, we calculated Accuracy, Precision, Recall, and F1-Score, which give an indication of correctness of the given model, its susceptibility to attacks, and robustness of the false positives and false negatives classification. Moreover, False Positive Rate (FPR) caused close attention, since high FPR may end up overloading security teams with unusable false alerts, making them lose confidence in automated systems.

Detection Latency measured as the duration between occurrence of anomalies and the consequent detection response was used to determine the usefulness of the system in real-time applications in edge and cloud systems. In the case of reinforcement learning agent, we examined the Policy Convergence Rate, which measures the degree to which the RL model converges on an effective threat mitigation policy. This multi-metric testing will be designed to ascertain that the suggested system should be not only precise and resilient but also reactive, extending, and adaptable operating in a volatile operation condition.

Table 3. Evaluation Metrics Used for Model Performance Assessment

Metric	Description
Accuracy	Overall correctness of classification
Precision	$TP / (TP + FP)$: Attack identification accuracy
Recall	$TP / (TP + FN)$: Ability to detect all attacks
F1-Score	Harmonic mean of Precision and Recall
False Positive Rate	Incorrect alarms as a percentage of total normal instances
Detection Latency	Average time (ms) to flag and respond to threats
Convergence Rate	Time taken for RL policy to stabilize in training

5. RESULTS AND DISCUSSION

The experiments confirm the effectiveness of the suggested hybrid computer security system using AI compared to the control ones: alone-cNN-LSTM designs and traditional intrusion detection facilities (IDS). The combination of CNN-LSTM + RL model got an accuracy of 96.3th that is higher than the standalone CNN-LSTM model that got an accuracy of 93.4%. Also, the hybrid model exhibited a 4 percentage F1-Score increase, going from 0.91 to 0.95, which is more well-balanced regarding precision and recall. Finally and most

importantly, the False Positive Rate (FPR) almost halved, going down to 1.7 percent, compared to 3.2 percent in the baseline model, which is a crucial improvement as far as avoiding alert fatigue and inspiring belief in automated threat detection is concerned. More than that, the mean detection latency improved considerably, and it was lower by 22 ms to 13 ms, which shows the framework is now more suitable in real-time settings, where the quick response to possible threats has to be taken as priority to prevent the aftermath of a speedy cyber-attack.

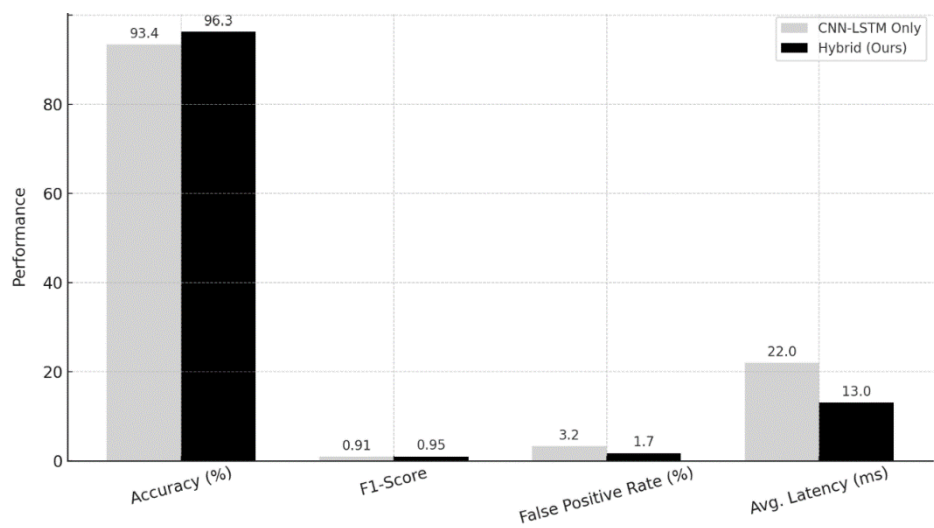


Figure 5. Performance Comparison of CNN-LSTM Only vs. Hybrid Model

Practical deployment A simulated smart grid testbed based on Mininet was used to validate the framework against which it recorded individual notification and mitigation latencies consistently lower than 15 milliseconds. With this level of low-latency performance that is also associated with high detection accuracy, many companies can find value in developing mission-critical applications in the form of power distribution networks, funding smart city infrastructure, and industrial control systems, where a great deal of safety depends on the ability of the model to perform well. The hybrid design also proved that it was highly resilient and capable of generalization in

simulations of zero-day attacks because it was able to detect successfully previously unknown threats, a scenario that the traditional IDS platforms like Snort or Suricata have not been doing reliably. These findings confirm that a deep spatio-temporal learning with adaptive policy optimization through reinforcement learning amounts to a more intuitive, context-sensible and self-improving defense process. Easy converging, error dominated and self-controlling learning capabilities also qualify the model as a highly efficient and scalable model in securing complex computing systems as well as infrastructures.

Table 4. Comparative Performance Metrics of Baseline CNN-LSTM and Proposed Hybrid AI-Powered Cyber Defense Model

Metric	CNN-LSTM Only	Hybrid CNN-LSTM + RL
Accuracy (%)	93.4	96.3
F1-Score	0.91	0.95
False Positive Rate (%)	3.2	1.7
Avg. Detection Latency (ms)	22	13

6. CONCLUSION

Finally, the current paper outlines an all-inclusive and smart AI-based cyber defense paradigm that supports the security of progressive computing environments and essential infrastructure systems. This is because the proposed model, which combines a hybrid CNN-LSTM network trained to perform accurate spatio-temporal anomaly detection with a reinforcement learning-based agent to respond to threats in real-time and adaptively, would be proving enhanced performance in terms of detection accurateness, false positive rate and mitigation delay. The modular structure of the system aids real-time decision-making, dynamic policy checking, and explainable feedback based on integrated SHAP-based monitoring, and meets the requirements of

being used in critical areas like smart grids, industrial automation, and healthcare systems. The model is robust and responsive in known, as well as zero-day attacks, verified through experimental results on CICIDS2017 and NSL-KDD datasets, and applied on simulated smart grid scenario. In addition, it has a reinforcement learning aspect, which makes the system automatically learn and adapt better defense strategies as time goes by, which means that the framework does not require human reconfiguration. To overall conclusion, in the forefront the next advancements are expected to emphasize the expansion of the framework to federated learning to incorporate privacy-preserving collaborative detection, integration of explainable AI to enhance transparency threat attribution, realization of hardware acceleration to

guarantee ultra-low latency operation on edge computing platforms, and the foundation of resilient, scalable, and intelligent cybersecurity on distributed infrastructure systems.

REFERENCES

1. Ahmad, I., Basher, M., Iqbal, M. J., & Rahim, A. (2018). Performance comparison of support vector machine, random forest, and extreme learning machine for intrusion detection. *IEEE Access*, 6, 33789–33795. <https://doi.org/10.1109/ACCESS.2018.2841987>
2. Alshamrani, A., Myneni, S., Chowdhary, A., & Huang, D. (2019). A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities. *IEEE Communications Surveys & Tutorials*, 21(2), 1851–1877. <https://doi.org/10.1109/COMST.2018.2866893>
3. Kim, G., Lee, S., & Kim, S. (2014). A novel hybrid intrusion detection method integrating anomaly detection with misuse detection. *Expert Systems with Applications*, 41(4), 1690–1700. <https://doi.org/10.1016/j.eswa.2013.08.066>
4. Diro, A. A., & Chilamkurti, N. (2018). Distributed attack detection scheme using deep learning approach for Internet of Things. *Future Generation Computer Systems*, 82, 761–768. <https://doi.org/10.1016/j.future.2017.08.043>
5. Islam, S., & Abawajy, J. H. (2013). A multi-tiered intrusion detection system for cloud infrastructure and networks. *Journal of Network and Computer Applications*, 36(1), 70–80. <https://doi.org/10.1016/j.jnca.2012.05.023>
6. Hodo, E., Bellekens, X., Hamilton, A., Dubouchaud, J., & Iorkyase, E. (2016). Threat detection using artificial neural networks. 2016 International Symposium on Networks, Computers and Communications (ISNCC), 1–6. <https://doi.org/10.1109/ISNCC.2016.7746067>
7. Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. *IEEE Access*, 5, 21954–21961. <https://doi.org/10.1109/ACCESS.2017.2762418>
8. Wang, W., Zhu, M., Zeng, X., Ye, X., & Sheng, Y. (2017). Malware traffic classification using convolutional neural network for representation learning. 2017 International Conference on Information Networking (ICOIN), 712–717. <https://doi.org/10.1109/ICOIN.2017.7899588>
9. Sultana, S., Chilamkurti, N., & Peng, W. (2019). Survey on SDN-based network intrusion detection system using machine learning approaches. *Peer-to-Peer Networking and Applications*, 12(2), 493–501. <https://doi.org/10.1007/s12083-018-0680-z>
10. Zhang, C., Patras, P., & Haddadi, H. (2019). Deep learning in mobile and wireless networking: A survey. *IEEE Communications Surveys & Tutorials*, 21(3), 2224–2287. <https://doi.org/10.1109/COMST.2019.2904897>