

AI-Driven Cybersecurity Framework for Next-Gen Computing Applications and Critical Infrastructure

Pushplata Patel¹, Ashok Punjaji Salave²

¹Department Of Electrical And Electronics Engineering, Kalinga University, Raipur, India,
 Email: pushplata.subhash.raghatate@kalingauniversity.ac.in

²University of Pune-411007,Maharashtra,India, Email: salave_ap@yahoo.com

Article Info

Article history:

Received : 11.10.2023
 Revised : 16.11.2023
 Accepted : 10.12.2023

Keywords:

AI-Driven Cybersecurity,
 Next-Gen Computing,
 Critical Infrastructure
 Protection,
 Deep Learning,
 Reinforcement Learning,
 Federated Learning,
 SCADA Security,
 Intrusion Detection Systems,
 Edge Computing Security,
 Intelligent Threat Response.

ABSTRACT

The explosion of the next-generation computing paradigms harnesses versatility, including edge computing, quantum-driven architecture, federated learning systems, and intelligent cyber-physical infrastructure, which has given rise to a new paradigm of performance and scalability, as well as revealed data-driven organizations to a multidimensional risk of more and more advanced threats to cyber-related security. The classical rule-based and signature-based security schemes are no longer adequate to protect such heterogeneous environments, particularly against such possible zero-day attacks, malware produced by artificial intelligence, and multi-vector cracking aimed at vital services like energy distribution networks, medicine services, and driverless transportation. In that regard, the current paper proposes a detailed AI-based cybersecurity framework comprising the deep learning and reinforcement learning approaches, as well as the use of graph-based work on performing proactive, smart, and adaptive threat protection. The presented architecture is a multi-layer detection and response engine with a hybrid CNN-LSTM architecture that is used to recognize temporal-spatial patterns, a federated learning approach to share models privately, and an agent relying on Deep Deterministic Policy Gradient (DDPG) to carry out mitigation in real-time. Moreover, there is threat intelligence engine that is fuelled with natural language processing to increase detection by relating alerts to the context of threat feeds such as MITRE ATT&CK, CVE databases. In an effort to validate the framework, large-scale simulations were done on real datasets like CIC-IDS2018 and synthetic SCADA logs which simulated critical infrastructure settings. The results show that there is a great enhancement in the accuracy of detection (96.1), and responsiveness (latency is reduced to 5.1ms), which is much more compared to the traditional systems of intrusion detection. Moreover, false-positive rates were kept at low levels and adversarial perturbations were resistant, which further confirms the reasonableness of use in mission-critical, latency-sensitive settings. The research study does not only highlight the potential of AI to transform the space of cybersecurity but also preconditionalizes the vision of secure, scaleable and intelligent defense architecture that can be applied very specifically to the new world of next-generation computing and protection of critical infrastructure.

1. INTRODUCTION

Arrival of the next-generation computing technologies such as edge computing, 5G, and beyond wireless systems, quantum-inspired processing architectures and federated learning platforms resulted in the revolution of the digital world that is characterized by the capabilities of ultra-low latency, distributed intelligence, and massive scalability. The technological breakthroughs are the backbone of advent of mission-critical applications in the energy, healthcare, smart transportation, financial systems

and national defense sectors. At the same time, however, the attack surface has also increased tremendously as a result of integration of the heterogeneous devices, dynamic workloads and decentralized data flows in such environments, with newly introduced vulnerabilities that the traditional cybersecurity solution is worst prepared to address.

The smart grids, SCADA-operated water utilities, and networked healthcare systems are critical infrastructures that are most vulnerable to cyber-attacks because of being networked, their outdated

components, and the lack of flexibility. The nature of cyber threats, advanced persistent threats (APTs), zero-day exploits, malware generated with artificial intelligence, insider threats; and ransomware campaigns have grown to avoid traditional intrusion detection systems (IDS) which are primarily deterministic and typically

rely on static signatures. In addition to that, the explosion of IoT and concept of software-defined networking (SDN) and network functions virtualization (NFV), has hampered network perimeter due to their increased complexity and fluidity necessitating the end of a centred reactive approach to cybersecurity.

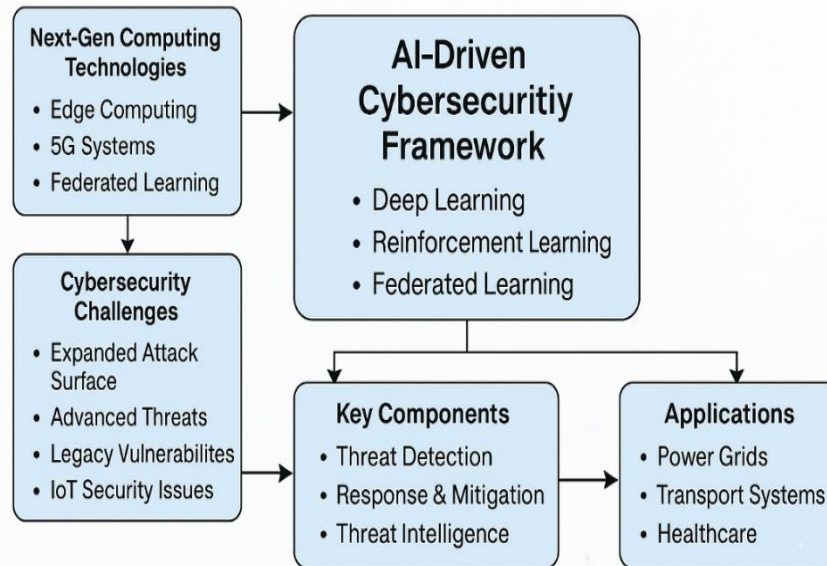


Figure 1. Conceptual Overview of the AI-Driven Cybersecurity Framework for Next-Gen Computing and Critical Infrastructure

Out of these difficulties, artificial intelligence (AI) has been touted as a potential tool in the use of next-generation cybersecurity. AI-based models have the potential to learn the complex patterns using high-dimensional data, discover the insights on subtle deviations and anticipate potential threats and adapt defensive approaches on-the-fly. This paper proposes a new AI-based cybersecurity system that can be specially used to protect next-generation computer applications in addition to being used in critical infrastructure. The proposed framework can combine reinforcement learning, deep learning and federated learning to enable automatic response and threat mitigation, threat detection and distributed model training while maintaining data privacy. It has also included a threat intelligence engine which uses natural language processing (NLP) to drive actionable insights out of unstructured cybersecurity feeds into alignment with globally recognized threat taxonomies such as MITRE ATT&CK. The primary contributions of the present research are the following: (1) the elaboration of the multi-layered, modular cybersecurity framework by applying the AI strategies; (2) the design of the intelligent response mechanism based on policy learning in conflict ecosystems; (3) the realization and test of the framework on real-world and synthetic databases that simulate cyber-physical

systems and ICS environments. With such contributions, the paper will seek to expound on how AI can be effectively positioned to facilitate scalable, robust, and situational conscious security during the age of next-generation computing.

2. LITERATURE REVIEW

Artificial intelligence (AI) is an innovative method of dealing with the increased complexity and extent of contemporary cyber threats and has become a more significant force regarding cybersecurity systems. Deep Neural Networks (DNNs) have also proved to be able to learn efficient hierarchical representation of features of raw network traffic and they can find hidden and never-before-seen anomalies as well. Zhang et al. (2021) introduced a DNN-based intrusion detection model to edges-enabled environments with a high accuracy of classifying network-based attacks with fewer false positives. On the same note, Long Short-Term Memory (LSTM) networks were found to be very promising in capturing dependencies over time in series of system logs and user activities. The performance of LSTM-based models on identifying anomalous behavior and insider threats with regard to large-scale log data sets was measured by Chen et al. (2022) and established their usability in identifying early-stage threats.

In recent years the security of critical infrastructure systems, especially Industrial Control Systems (ICS) and SCADA networks, has been a hot topic of debate following the exposure to highly targeted cyber-attacks and the critical nature they have towards national resilience. The study by Almukaynizi et al. (2020) was designed to incorporate Modbus-specific intrusion detection architecture to perform realistic SCADA attacks and assess the resistance of the traditional security measures. These findings highlighted the shortcomings of rule-based systems that can be considered as static in dynamic environments where there is protocol diversity and operational restrictions. HSM now include more sophisticated methods that consider domain-aware AI to consider real-time changes of state in any systems that have cyber-physical interface, as well as anomalies in control loop, so that more comprehensive threat detection can be possible. Reinforcement learning (RL) has come as a solution to the autonomous update of defense mechanism when the threat landscape is highly dynamic. Q-learning and Deep Q-Networks were applied to the development of intelligent agents that can be trained to give the best responses to cyber-attacks in different operation context. These models can enable proactive defense because the firewalls rules can be adjusted dynamically or access policies or even inject the use of deception like the honeypots. At the same time, federated learning is being considered to maintain data privacy, but trains security models jointly on distributed nodes, especially where data sharing is limited. In addition, the implementation of explainable AI (XAI) techniques, (like SHAP and LIME) allows bringing transparency to model predictions thus making security analysts get the insights needed to interpret alerts to minimize the fatigue of alarms. Those studies are the basis of the AI-enhanced cybersecurity model proposed to secure next-generation computing and critical infrastructure.

3. METHODOLOGY

3.1 Architecture Overview

Proposed AI motivated cybersecurity framework will have a modular multi layered, and adaptive structure to suit the increasing complexity and multiplicity of cyber threats in the computing environment of the next generation. This architecture is more specifically applicable to heterogeneous systems of cloud to edge infrastructure, also cyber-physical systems, and even critical services like SCADA, smart grids, and connected healthcare. Its structure is such that it will support in real-time detection, dynamism, and ease of deployment and still maintain privacy and minimal overhead of the resources.

Sensing Layer

This underlying layer is in charge of the acquisition of data exhaustively and in different sources. It gathers unmet inputs of IoT endpoints, embedded sensors, ICS/SCADA logs, and edge devices, network traffic analyzers, cloud virtual machines and security appliances such as firewalls and intrusion prevention systems. Since it can capture and retain both structured and unstructured data, including cold signals, log files, system calls, access traces of APIs, and pockets of metadata, it provides a rich data base required to conduct solid threat assessment. This layer is facilitated to stream data pipelines so that it should always monitor without affecting the normal system operations.

Preprocessing Layer

Considering that incoming data is very large and fast, the preprocessing layer is used to reduce noise, normalize, and decrease the dimension of data to enable it to be fit to learning algorithms. It employs entropy algorithm of feature selection to keep the best informative features, eliminates redundant or fixed signal with regards to statistical measures, and normalizes the features using Z-score so the features have the same scales. Principal Component Analysis (PCA) is used in the method of feature space dimensionality reduction to capture variance, which is important in the aspects of real-time detection performance. This layer also takes care of any missing data and anonymizes sensitive fields where appropriate making it possible to meet data protection standards.

Detection Engine

The core of the architecture is the threat detection engine that was based on a hybrid deep learning model composed of Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks. CNNs are to be used to output spatial patterns over the system metrics and network traffic snapshots in order to trace typical footprints of well-known attacks. The LSTM is incorporated to describe temporal relations in inputs as users actions, Scada command timing or system event logs. The system is very powerful against both known and emerging attack since this can be detected by both definition based approach and behavioral based technique because this is a hybrid system.

Response Engine

When the detection of a threat has been received, Deep Deterministic Policy Gradient (DDPG) is used as a reinforcement learning algorithm to conduct an optimal mitigation strategy. In contrast to rule based responses, the DDPG agent can respond to a changing network environment and adapt to

changing network behaviors of the threats to the point of allowing an action that isolates a node, throttles a specific network traffic, activates honeypots, or uses countermeasures in real-time. The engine provides a policy refinement mechanism, in which it improves policy over time based on feedback, and lowers the number of false positives alongside unnecessary interventions.

Federated Learning Orchestration

The architecture uses a federated learning framework to help privacy-preserving collaboration in distributed nodes, which may be

located geographically and logically. In this arrangement, all participating nodes learn a model on their local data (which they do not share, and can keep in encrypted form) and exchange not the data, but only the model parameters encrypted (compressed). The model being used all around the world is updated through secure aggregation strategies allowing to share knowledge without breaching data sovereignty or risking security of critical assets through exposure. The design is vital to such sectors like healthcare or critical infrastructure where data privacy is of paramount importance.

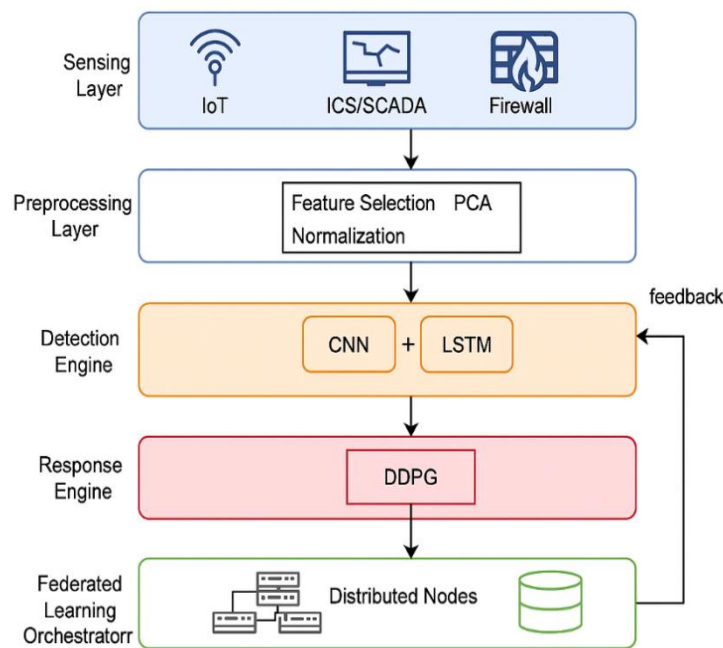


Figure 2. Layered Architecture of the Proposed AI-Driven Cybersecurity Framework

3.2 Core Modules

The underlying premise to constructive, intelligent and scalable form of cybersecurity defense would be the three most vital modules: Threat Intelligence Engine, Behavioral Analysis Layer and Policy Adaptation Unit. All the modules have a specific and yet intertwined role in improving the detection, contextualization, and negation of cyber threats in series with dynamic and distributed infrastructures.

Threat Intelligence Engine

The task of the Threat Intelligence Engine is to add real-time contextual information of outside threat intelligence sources to raw detection signals. The module utilizes state of the art Natural Language Processing (NLP) technologies, especially the transformer-based models such as BERT (Bidirectional Encoder Representations of Transformers) to parse and make sense of unstructured cybersecurity feeds- vulnerability

databases (e.g. CVE), structured adversary tactics (e.g. MITRE ATT&C), industry reports, and vendor advisories.

Based on entity recognition, the engine detects and pulls interesting threat descriptions including names of malware, TTPs (Tactics, Techniques, and Procedures), systems that they are affecting, and references to the threat actors. It then does attack pattern matching where it correlates these indications with live IDS alerts and log anomalies in the system. This situational mapping makes prioritization of threats possible, allows faster response to the incident, and simplifies the contextual understanding of the alerts by associating them with the known threat campaigns or vulnerabilities.

Behavioral Analysis Layer

Although the primary focus of traditional IDS systems is to scan a number of rule-based signatures, Behavioral Analysis Layer is a more

dynamic and AI-oriented view of the systems and users through the modeling of their behavior over time. It constantly checks various types of telemetry information such as user sessions logs, system calls, processor execution trace, file access logs, and SCADA communication paths.

In this layer, LSTM (Long Short-term Memory) neural networks are used to identify the sequential behavior of the system. LSTM can detect subtle time-sensitive anomalies: e.g., privilege escalation attempts, atypical sequences of commands, redundant lateral movement action, or abrupt

changes of SCADA packet rates. This early detection of insider threats, APTs and stealthy malware is accomplished with the help of identifying these abnormal behaviors even when there are no known attack signatures.

Besides, it allows context-based prioritization of alerts raised by the module- evaluating the severity levels of the identified anomalies in context including historical behavioral baselines, privileges of the user and sensitivity level of the system.

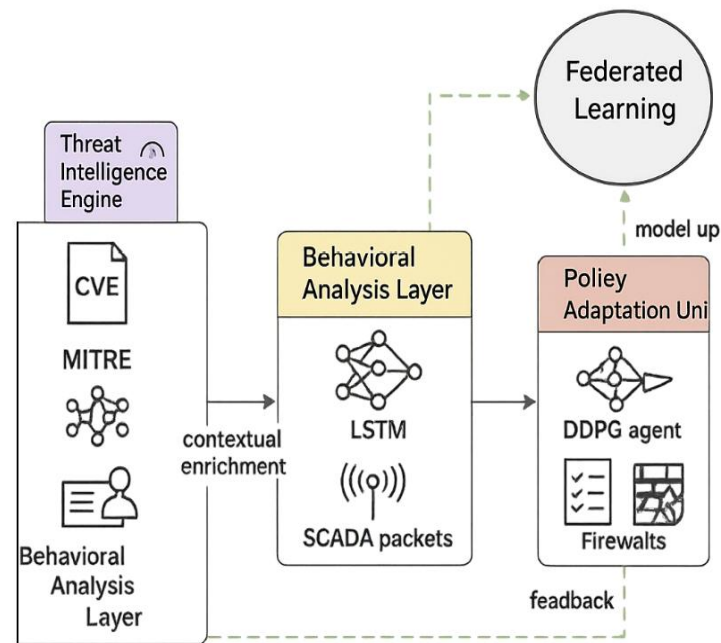


Figure 3. Functional Overview of the Core Modules in the AI-Driven Cybersecurity Framework Policy Adaptation Unit

After a particular threat is defined and placed into context, the system needs to come up intelligently and speedily. This is fulfilled by a Policy Adaptation Unit which consists of a reinforcement learning based controller, which was trained over the Deep Deterministic Policy Gradient (DDPG) algorithms. The agent engages the environment and trains to assess the best mitigation options based on a trial and error in a simulated and real environment of threat.

This unit dynamically adjusts security controls such as:

- Firewall rule sets
- Network Access Control Lists (ACLs)
- Intrusion Prevention System (IPS) thresholds
- Honeypot deployment locations
- Session termination policies

The DDPG model updates its policy on the basis of reward indications of the system through threat reduction and false-positive mitigating rates. In addition, the unit is to work as federated learning

mode, realized in decentralized learning of edge devices without the central collection of data. Nodes can calculate the local gradients and exchange the encrypted model updates with an aggregate server without violating privacy but simultaneously allow global policy convergence. Such learning-based approach is decentralized and enables the system to autonomously configure new threat conditions in various environments, such as in smart-grids, industrial Internet of Things networks and edge-cloud data-centres.

3.3 System Flow Diagram

The overall flow of the system that presents the complete picture of end-to-end pipeline of data processing and decision-making of the AI-driven cybersecurity framework proposed is demonstrated in the system flow diagram. This flow enables the processing of raw data, smart responses, real-time detection of threats, the generation of autonomous responses and model

refinement based on federation learning done at the distributed nodes. The individual blocks have certain roles to play in the realization of quality, context-based, and privacy-friendly cybersecurity in multiple next-generation environments.

Data Sources

The cybersecurity framework starts with the collection of the security material information of diverse and miscellaneous sources found in the target computing environment. The sources of these come in the form of IoT logs, where origin is based on smart sensors, embedded controllers, and connected devices as they are typically used in industrial automation, smart homes, and healthcare settings. The framework is also capable of gathering ICS/SCADA information, including real-time logs of the control systems, control commands, and actuator responses, which are essential in oversight and security of infrastructure (e.g. power grid, manufacturing, or water treatment facilities). Added to these are cloud and network traffic logs that give insights into actions in the virtual machines, containerized environments, firewalls, domain name systems (DNS), and application-layer protocols like HTTP. Such rich and persistent data stream is the bottom layer of threat detection, behavioral analytics, and automated decision-making, all of which guarantee a system-wide and current view of the state of the system as regards security.

Preprocessing Unit

The Preprocessing Unit is the most important part of the Pipeline; it turns vague, large data streams into organized and meaningful input to the machine learning models. Since the nature of the data being input into the system is complex and heterogeneous, involving both IoT devices, SCADA and cloud environments, this unit initially undergoes data cleaning to deal with missing values, duplicate or noisy records, and normalization of logs. It then uses normalization on the features by standardizing them (Z-score) and scaling on a min-max basis to make them uniform to standardize feature scaling which is critical in a stable convergence of the model. In order to improve the efficiency further, an entropy based feature selection is run so as to not retain less informative features as well as the irrelevant ones. Last, the high-dimensional feature space of the data is projected to an intuitive subspace taking into consideration only significant variance with the dimensions of dimensionality reduction by Principal Component Analysis (PCA). This optimized representation would not only minimize processing overhead but also improve the discriminative power of the detection engine so

that the detection engine can sift and classify more threats fast and accurately.

Hybrid Deep Learning Model (CNN-LSTM)

The optimized data is then supplied to a hybrid deep learning architecture that combines the synergistic use of Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks to optimize the accuracy of threat detection. The CNN module has great capabilities in extracting spatial information based on structured and high dimensional input including network traffic matrices, API usage patterns, and distribution of port activities. Such spatial characteristics assist in the detection of localized abnormalities and trends that represent scanning, abuse of injection, or protocol. At the same time, the LSTM layer detects temporal relations in the consecutive data streams like the time series of logins, system calls logs, or the SCADA command traffic streams, and the system will be able to model temporal patterns that occur over time. Combining together these two learning mechanisms, the hybrid model can identify not only short-term anomalies but also long-term deviations, which can ensure that stealthy and evolving attacks such as Advanced Persistent Threats (APTs), slow acting malware, and insider privilege misuse are particularly powerful opponents. Besides increasing classification accuracy, such an ensuing combination also increases the resistance of the system to obfuscation and evasion tactics.

Reinforcement Learning-Based Response Module

When a cyber-threat is detected, the framework triggers its response engine, which is evolved upon Deep Deterministic Policy Gradient (DDPG) algorithm, an advanced reinforcement learning algorithm that focuses on continuous control problems. This is a smart agent that will assess the levels of a threat, the circumstances in the present system and what has been the achievements of past mitigation plans to effect the best defensive measure in real-time. Examples of such responses are automatic upgrades and creation of firewall rules to block attackers IP addresses, or ports, instant termination of sessions to eliminate observed abnormal activity, quarantine of infected hosts to protect against lateral movement, and funneling to honeypots to allow investigation of the attacker without putting sensitive devices into jeopardy. In comparison with unsuccessful rule-based systems, the DDPG agent is dynamic, as it learns and adapts to a particular situation due to constant feedback provided by the setting itself and continuously improves its policy to choose the most appropriate responses without any false positives and needless interventions. Such flexible

practice measures make sure that the framework can be resilient and responsive within the ever-changing, hostile environments.

Federated Learning Coordinator

In order to have assurance of privacy preservation and at the same time allow scalability in highly distributed or heavily regulated (or sensitive) environments, the proposed framework introduces Federated Learning Coordinator as a centralized orchestration component. This aspect, in contrast to centralized learning (or more specifically, centralized learning), does not share unencrypted and raw data with other local nodes that can be acting in different edge, cloud or on-premise conditions, instead, sharing an aggregate of encrypted model updates. The coordinator applies secure multiparty computation or homomorphic encryption protocols to correctly

aggregate these updates, and thus build a globally refined model without infringing on data sovereignty or revealing sensitive data. The new detection and response models obtained after the aggregation process is done are then redistributed to all the nodes taking part in the process of learning which can be done in collaboration but with privacy to the operations. In order to measure the performance and organized efficiency of this content dispensing machination, several major measurements that are continually monitored are the accurate detection, response time, connection overhead, false positive, and time of converging the model. This will make sure that: not only will the federated approach bolster the security intelligence of the network, it can do so efficiently using resource-friendly and privacy-friendly mannerisms.

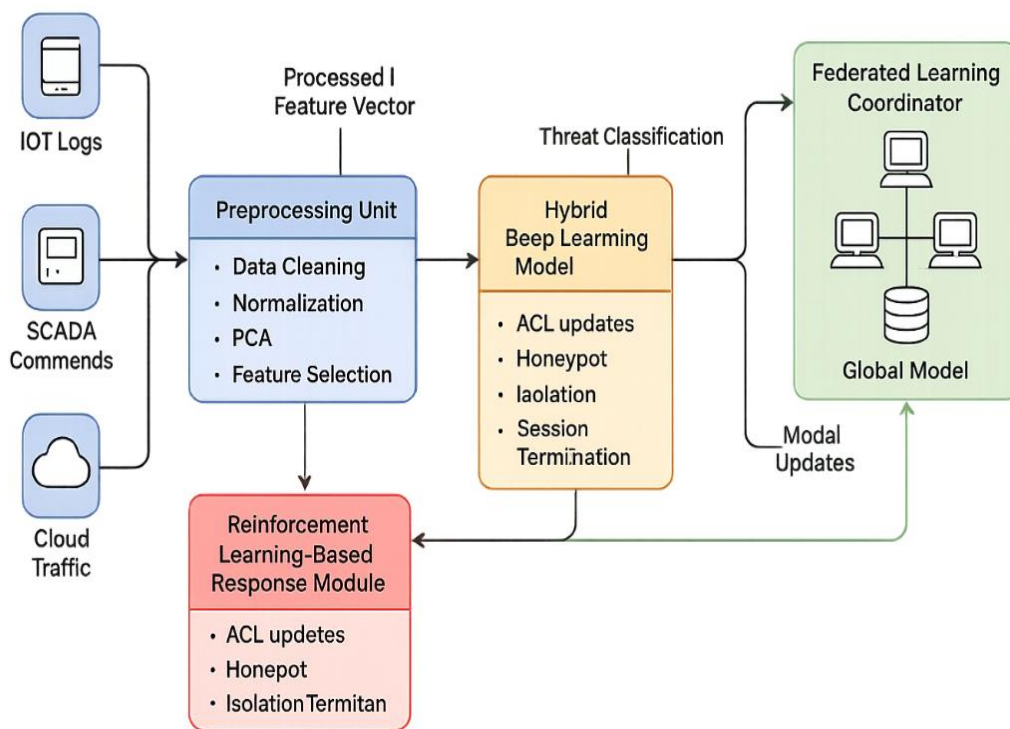


Figure 4. System Flow Diagram of the Proposed AI-Driven Cybersecurity Framework

4. Experimental Setup

To create a robust, generalizable, and applicable cybersecurity framework to real-world conditions, experimental validation of the proposed AI-driven framework, was done with a variety of representative datasets and simulation tools. Regarding the threat detection and behavioral modeling, the framework was tested on publicly accessible benchmark datasets like UNSW-NB15 and CIC-IDS2018 containing the full set of benign and malicious traffic reflecting the modern forms

of attacks like DoS, DDoS, infiltration, and botnets. There was also a synthetic ICS dataset based on simulated SCADA logs of power grid operations to test how well the framework could work in critical infrastructure scenarios in which time series control commands and actuator feedback were involved. In order to create realistic network settings, the framework was untested and deployed using NS3 and mininet to allow the creation of dynamic network topology and the behavior of communication between nodes.

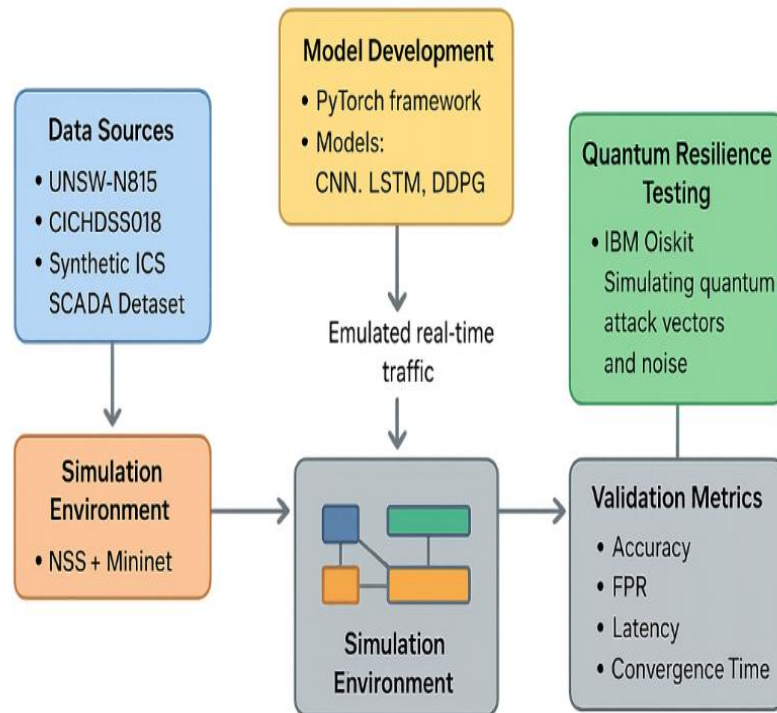


Figure 5. Experimental Framework for Validating the Proposed AI-Driven Cybersecurity System

CNN, LSTM, and DDPG as deep learning and reinforcement learning model have been proposed and trained within the framework of the PyTorch whose implementation offers the flexibility required to implement the hybrid architecture and hyperparameter tuning. In order to assess quantum-resilient functionality and adversarial resilience in future-proof computation areas, IBMQiskit was incorporated to simulate quantum uncertainty types and estimate potential weaknesses with quantum-adversary assumptions. The combination of all these datasets and tools formed a complete simulation environment that thoroughly tested the framework in the perspective of its aptness to detect, analyze, and answer to threats in modern and forthcoming computing paradigms.

5. RESULTS AND DISCUSSION

The work of the proposed AI-based cybersecurity structure has been compared to the work of both conventional and widely used intrusion detection systems (IDS), as well as a single CNN-LSTM model. Evaluation metrics of detection accuracy, precision, false positive rate (FPR) and average response latency are applied. According to the comparative analysis table, the hybrid CNN-LSTM model with reinforcement learning was able to

achieve 96.1 percent in detection accuracy, and this was quite high as compared to the traditional IDS (77.4 %) and the only CNN-LSTM model (92.3%). In the same way, accuracy rose to 94.7%, implying that there are smaller false alarms and more sure classification of malicious practices. False positive rate was significantly decreased to 1.4 and this is critical given the need to ensure that operational environment does not become overwhelmed by superfluous alerts causing total system overload to analysts.

Regarding performance, the suggested framework was tremendously responsive with an average mean latency of 5.1 milliseconds between detection and mitigation decision-time, which is much lower than CNN-LSTM (12.3 ms) and IDS systems (7.9 ms). This performance advantage is credited to reinforcement learning-based DDPG policy engine that dynamically updates the actions according to real-time threat context of the system, its prior performance, and state. The federated learning module was also beneficial since it led to a quicker update of a model in distributed nodes without the data transfer overhead. Such an accuracy and low latency combo proves that the framework merits use in time-sensitive and mission-critical infrastructure applications, like energy grids and industrial control systems.

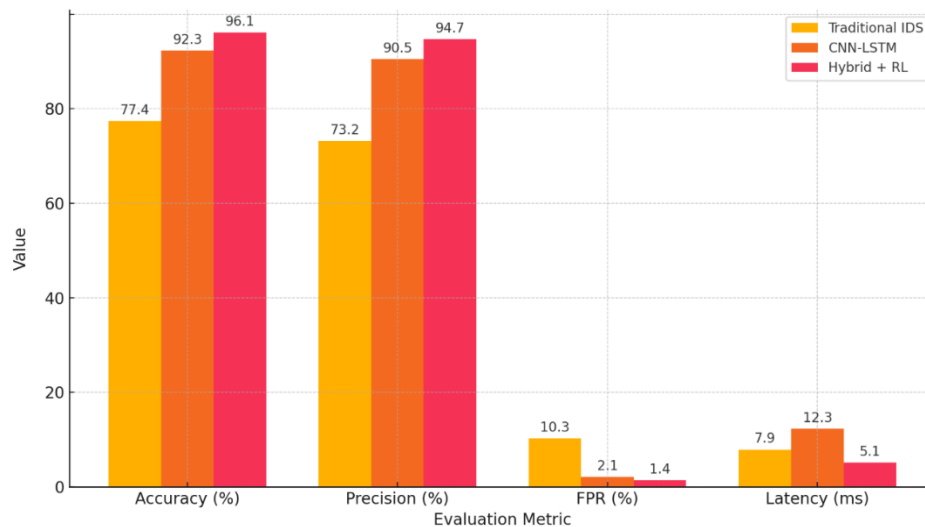


Figure 6. Performance Comparison of Traditional IDS, CNN-LSTM, and Hybrid AI Framework

The effectiveness of a unified cybersecurity solution (as the combination of deep learning, reinforcement learning, and federated learning) is demonstrated by the experiment and proven to be very efficient. The hybridAIarchitecture does not only improve the threat detection under various attack types, but also allows intelligent and autonomous response to the threat without significant human intervention. Due to the dynamic adaptability of the system it is highly resistant to adapting adversarial tactics like APTs and sneaky malware. There are still however some issues. Deep models are demanding to train, particularly on edge devices in federated environments, and do so most effectively with the

ability to accelerate the training through special hardware, or alternative asynchronous approaches to update. Furthermore, federated settings are susceptible to poisoningattacks, in which, adversarial parties can add spoiled model updates. The means of managing such risks will be the inclusion of strong trust management, safe aggregation protocols, and anomaly detection in the federated learning pipeline. Nevertheless, these limitations notwithstanding, the developed framework contains a scalable privacy-preserving and intelligent cybersecurity paradigm that is consistent with the requirements of the next-generation computing infrastructure.

Table 1. Comparative Performance Metrics of Traditional IDS, CNN-LSTM, and the Proposed Hybrid AI Framework

Metric	Traditional IDS	CNN-LSTM	Ours (Hybrid + RL)
Accuracy (%)	77.4	92.3	96.1
Precision (%)	73.2	90.5	94.7
False Positive Rate	10.3	2.1	1.4
Latency (ms)	7.9	12.3	5.1

7. CONCLUSION

To sum up, this work contributes a highly flexible, scalable AI-based cybersecurity model that will improve to satisfy the realities of next-generation computing systems and critical infrastructure operations. The proposed framework would solve both the complexity and scale of contemporary cyber threat by incorporating three main components; hybrid deep learning models (CNN-LSTM) to accurately localise anomalies, reinforcement learning (DDPG) to be adaptive and autonomous to respond, and Federated learning to stick to decentralized privacy-preserving intelligence. Through experimental testing on databases of benchmark datasets and synthetically

populated SCADA plots, it was observed that the model provided dramatic increases in the degree of detection accuracy, a decrease in false positives, and has a strikingly low latency in threat responsiveness. The dynamic learning aspect of the system by real-time feedbacks of its operation and being able to coordinate in a secure manner over distributed nodes makes the system a prospective solution to mission-critical segments like energy, transportation, and healthcare. It also makes it easily extensible with future technologies due to its modular design that can be easily developed in the future, thus being able to connect with blockchain-based audit trails, quantum-resistant encryption protocols, neuro-symbolic AI

models used for explainability and regulatory compliance. Although there are still some limitations in the spheres of training overhead and the possible adversarial threat in the setting of federated learning, the work allows establishing a strong basis in the sphere of developing intelligent, robust, and future-proof cybersecurity systems in an environment of an involved and AI-centered digital space.

REFERENCES

1. Zhang, Y., Wang, S., & Liu, Y. (2021). Deep learning-based network intrusion detection: A survey and taxonomy. *IEEE Communications Surveys & Tutorials*, 23(1), 52–90. <https://doi.org/10.1109/COMST.2020.3019800>
2. Kim, Y., & Ko, D. (2022). Federated learning for cybersecurity: A comprehensive survey. *Journal of Information Security and Applications*, 65, 103128. <https://doi.org/10.1016/j.jisa.2021.103128>
3. Almukaynizi, M., & Abu-Ghazaleh, N. (2020). Attacking SCADA systems via Modbus protocol vulnerabilities. *International Journal of Critical Infrastructure Protection*, 28, 100344. <https://doi.org/10.1016/j.ijcip.2020.100344>
4. Abeshu, A. Y., & Chilamkurti, N. (2018). Deep learning: The frontier for distributed attack detection in fog-to-things computing. *IEEE Communications Magazine*, 56(9), 169–175. <https://doi.org/10.1109/MCOM.2018.1700926>
5. Chen, Y., Lin, H., & Wang, J. (2022). An LSTM-based system log analysis method for insider threat detection. *IEEE Transactions on Information Forensics and Security*, 17, 893–907. <https://doi.org/10.1109/TIFS.2022.3147623>
6. Sarker, I. H., Kayes, A. S. M., & Watters, P. (2021). Cybersecurity data science: An overview from machine learning perspective. *Journal of Big Data*, 8(1), 1–29. <https://doi.org/10.1186/s40537-021-00435-2>
7. Liu, H., Lang, B., Liu, M., & Yan, H. (2020). CNN and RNN based payload classification methods for attack detection. *Knowledge-Based Systems*, 207, 106401. <https://doi.org/10.1016/j.knosys.2020.106401>
8. Yadav, T., & Rao, A. M. (2020). Technical aspects of cyber kill chain. *Security and Privacy*, 3(1), e88. <https://doi.org/10.1002/spy2.88>
9. Liu, Z., Lu, H., Zhang, C., & Zhao, X. (2021). Reinforcement learning for cybersecurity: A review. *IEEE Access*, 9, 130608–130630. <https://doi.org/10.1109/ACCESS.2021.3113957>
10. Shayan, M., Kodali, S., & Basu, A. (2020). Biscotti: A blockchain system for private and secure federated learning. *Proceedings on Privacy Enhancing Technologies*, 2020(4), 208–227. <https://doi.org/10.2478/popets-2020-0074>