

AI-Based Threat Monitoring Framework for Critical Infrastructure in Developing Economies

Emmanuel Eturpa Salami, Precious Bamiyo Metiboba, Caleb Lateef Umoru, Tajudeen Isah, Vincent Onuwabagbe Ogbeide

Confluence University of Science and Technology, Kogi State

DOI: <https://dx.doi.org/10.51584/IJRIAS.2026.110400157>

Received: 14 April 2026; Accepted: 19 April 2026; Published: 16 May 2026

ABSTRACT

The increasing digitalization of critical infrastructure by developing countries has opened new security problems that existing conventional methods have failed to address. Particularly with advanced threats rapidly evolving. This paper aims to develop a system that monitors both active and passive threats using Artificial Intelligence integrated with lightweight deep learning for optimization to watch for threats in places where resources are limited. We used an hybrid model that combines Autoencoder and LSTM. The model performed excellently in learning threats. We used 20 epochs in training the system and observed a stable convergence when training the epochs and when it was tested, which means it performed good in terms of its generalization capacity. The reconstruction error distributions result showed a significant separation between benign and anomalous events ($p < 0.01$). Our tests showed that the proposed threat model framework was very good at detecting threats with accuracy of 96% and a confidence range of 94.95% to 97.05%. We did a two-sample t-test comparing reconstruction errors between normal and attack traffic that produced a statistically significant separation ($t \approx 18.7$, $p < 0.001$) with threat detection score of 95%. This showed the model performed well at telling the difference between attack traffic. All these results together show that our model is reliable and can work well in places where resources are limited. The outcome is a system that can watch for threats and adapt to situations, within countries that are resource constrained.

(Keywords: Artificial Intelligence, Threat Monitoring, Critical Infrastructure, Developing Economies, Cybersecurity, Anomaly Detection)

INTRODUCTION

Critical infrastructure like energy grids, water supply systems, telecommunications and transport networks is very important for the stability of a country. In developing economies these systems are becoming more digital. They are not secure because there are not enough people with the right skills, enough capital and good defense technologies (Madan & Bedi, 2022). When these systems are attacked it can disrupt services and cause economic losses and safety problems (Alrashdi et al., 2020). Artificial Intelligence (AI) provides a promising paradigm for proactive cybersecurity. The ability of AI to learn, adapt, and detect complex threat patterns unseen by existing conventional systems enables a continuous monitoring and early threat anticipation (Lin et al., 2021). Nonetheless, applying AI-based security by developing countries presents unique challenges such as inconsistent data availability, computational limitations, and infrastructural heterogeneity (Ahmed et al., 2016).

Despite advances in AI-based cybersecurity, developing nations continue to rely on reactive, signature-based detection systems incapable of identifying zero-day or polymorphic attacks. Such limitations leave critical infrastructures vulnerable to ransomware, insider threats, and targeted sabotage (Ahmed et al., 2016). Major challenges identified to be existing in this regard are: Limited computational resources, as most AI frameworks require high processing power, unavailable in typical infrastructure management centers. Also, the absence of localized, labeled datasets reflecting regional attack patterns hinders model training and generalization (Amin et al., 2021), and disconnected infrastructure sectors prevent collaborative threat intelligence sharing (Ogundikun et al., 2023). Hence, there is a pressing need for an AI-driven, resource-optimized monitoring solution capable of detecting and mitigating threats in real time while adapting to the unique operational realities of developing

economies. This study therefore, proposes an AI-based threat monitoring framework specifically tailored to the context of developing economies. The goal is to design a lightweight, scalable, and intelligent system capable of detecting real-time cyber anomalies within critical infrastructure networks without relying on high-performance computing or expensive proprietary tools.

The intersection of AI and cybersecurity has evolved significantly in the last decade. Research has demonstrated that AI-driven intrusion detection systems outperform traditional rule-based models in adaptability and detection accuracy (Aldweesh et al., 2020). For instance, deep learning architectures particularly convolutional neural networks (CNN) and long short-term memory (LSTM) networks are widely applied for anomaly detection due to their ability to capture spatial-temporal correlations in network traffic (Kim et al., 2022). Alrashdi et al., (2020) presented an IoT-based intrusion detection framework using machine learning algorithms, achieving substantial improvements in attack recognition over signature-based methods. Similarly, (Lin et al., 2021) developed a lightweight deep learning model for Industrial IoT systems, demonstrating that compressed networks can achieve high detection rates while maintaining computational efficiency. These studies confirm the viability of AI for cyber defense in real-time environments.

However, most of the existing frameworks assume the availability of high-quality datasets and computational infrastructure, conditions rarely met in developing economies. A key gap in research lies in designing an AI models that are both resource-aware and context-sensitive capable of functioning effectively where internet connectivity, processing power, and skilled personnel are limited (Ogundikun et al., 2023). Furthermore, recent advances in federated learning and edge-based AI have shown promise for decentralized cybersecurity monitoring, allowing distributed devices to collaboratively detect threats without centralized data aggregation (Zhang & Xu, 2022). Yet, few studies have contextualized these innovations for critical infrastructure protection in emerging economies.

Conceptually, socio-technical systems (STS) frameworks emphasize that cybersecurity for critical infrastructure is not purely technical; it involves interactions among people, processes, and technology. STS-based security modelling helps integrate organizational practices, human factors, and technical controls to create robust security postures (Ani et al., 2023; Madan & Bedi, 2022). While existing theoretical frameworks commonly applied to critical infrastructure security include resilience engineering, defense-in-depth, and the observe–orient–decide–Act (OODA) loop for rapid decision-making. Resilience frameworks focus on system capacity to anticipate, withstand, recover, and adapt to adverse cyber events (Estay et al., 2020). defense-in-depth prescribes layered security controls across physical, network, application, and human domains to reduce attack surface and provide fallbacks. The OODA loop has been adapted for cyber defense to improve situational awareness, accelerate response cycles, and support continuous learning in security operations (Husák et al., 2022).

Existing Ai-Based Frameworks For Threat Monitoring Of Critical Infrastructure

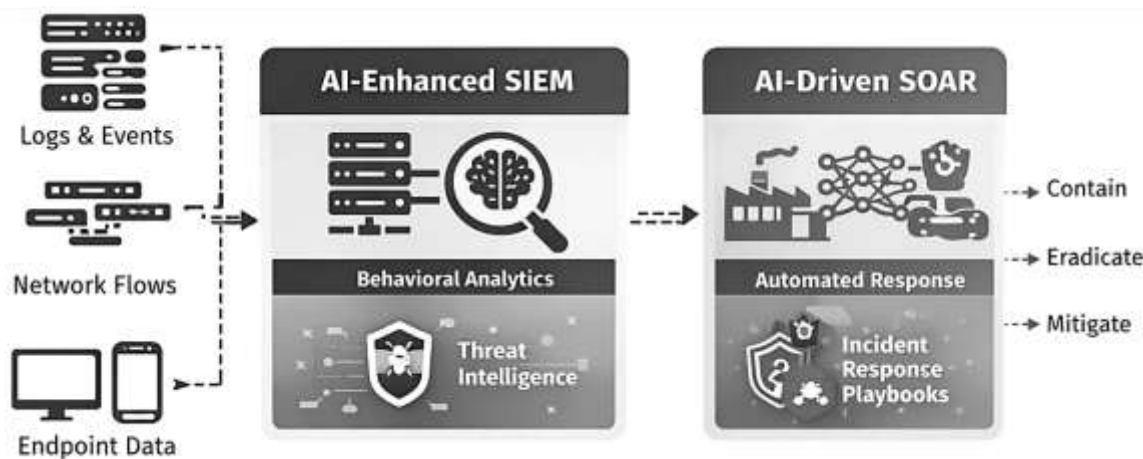


Figure 1: AI-Enhanced Security Information and Event Management (SEIM) & Security Orchestrated Automation

Figure 1 shows the AI-enhanced security information and event management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms which represent one of the most widely adopted AI-based threat monitoring approaches for critical infrastructure globally. These systems employ machine learning models to correlate logs, network flows, and endpoint telemetry across large-scale environments to identify advanced persistent threats (APTs), insider attacks, and zero-day exploits (madan & Bedi, 2022). Modern SIEM platforms integrate supervised and semi-supervised learning for behavioral analytics, anomaly detection, and threat scoring. For example, IBM QRadar Advisor leverages natural language processing and knowledge graphs to enrich alerts using threat intelligence feeds, while Splunk applies user and entity behavior analytics (UEBA) to detect deviations from normal operational baselines (Husák et al., 2022).

Despite their effectiveness in well-resourced environments, these frameworks exhibit significant limitations. High Computational and storage overhead necessitate the system to rely on continuous ingestion and centralized processing of massive volumes of logs and telemetry data, which requires high-performance servers, large-scale storage, and reliable high-bandwidth connectivity resources often unavailable in developing economies (Ogundokun et al., 2023). Dependence on Skilled Cybersecurity Personnel

is another known limitation. Effective operation requires trained security analysts capable of tuning models, interpreting alerts, and managing automated playbooks. In many developing nations, shortages of cybersecurity professionals severely limit operational sustainability [4]. The problem of Licensing and Vendor lock-In costs makes commercial AI-SIEM platforms very costly, incurring substantial licensing, maintenance, and subscription costs, which makes long-term adoption economically infeasible for public-sector critical infrastructure operators in low-income regions (madan & Bedi, 2022) Limited ICS/SCADA Semantic Awareness because most SIEM solutions are originally designed for enterprise IT environments. Their AI models lack deep contextual understanding of industrial control system (ICS) protocols, leading to false positives or missed attacks in operational technology (OT) networks (Zhao et al., 2022).

Developing economies typically operate fragmented, aging infrastructure with intermittent connectivity and limited budgets. Centralized AI-SIEM architectures amplify single points of failure and struggle under unreliable power and network conditions. Moreover, the lack of localized threat intelligence datasets reduces model relevance, as most platforms are trained on attack patterns from developed regions (Ogundokun et al., 2023).

Deep Learning-Based Intrusion Detection Frameworks for Industrial Control Systems

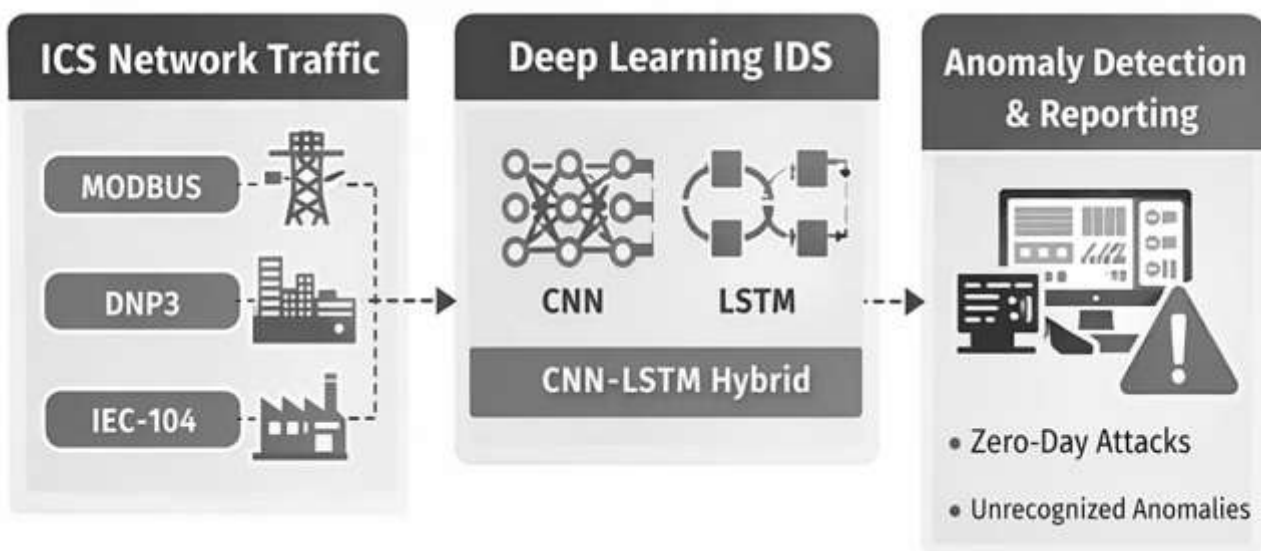


Figure 2: Deep Learning-Based Intrusion Detection Frameworks for Industrial Control Systems

Another prominent class of AI-based threat monitoring frameworks involves deep learning intrusion detection systems (IDS), as shown in Figure 2, which is tailored for industrial control systems and smart infrastructure. These frameworks commonly employ Convolutional Neural Networks (CNN), long short-term Memory (LSTM)

networks, or hybrid CNN-LSTM models to capture spatial and temporal dependencies in network traffic (Kim et al., 2022). Such frameworks have demonstrated high detection accuracy against known and unknown attacks in controlled experimental environments. For instance, CNN-LSTM models effectively detect command injection and replay attacks by learning temporal patterns in Modbus, DNP3, and IEC-104 traffic (Lin et al., 2021).

While academically robust, these frameworks present several practical limitations. The issues of Data Dependency and Labeling Challenges make most deep learning IDS frameworks require large volumes of labeled attack data for training. In developing economies, such datasets are scarce due to limited incident reporting and a lack of mature threat intelligence pipelines (Amin et al., 2021). Also, Model Complexity and Resource Intensity makes CNN-LSTM architectures computationally expensive, making real-time deployment difficult on legacy ICS hardware commonly found in developing nations (Ani et al., 2023). Centralized Training Assumptions is what make many frameworks assume centralized data aggregation for model training, which introduces privacy risks, bandwidth bottlenecks, and single points of compromise (Zhang & Xu, 2022). The issue of Operational Fragility makes these models to lack resilience mechanisms. Model retraining failures, concept drift, or communication outages can degrade detection performance without graceful degradation [6]. The challenge of Non-Adoptability in Developing Economies. The operational realities of developing economies, unreliable power supply, limited compute infrastructure, heterogeneous devices, and minimal cybersecurity staffing make large deep learning IDS deployments impractical. Consequently, many critical infrastructure operators default to static rule-based systems despite their known weaknesses.

This gap motivates this paper’s contribution: an AI-based threat monitoring system explicitly designed for low-resource, high-risk contexts typical of developing nations’ infrastructure ecosystems. Integrating AI techniques with socio-technical and resilience-based theoretical models suggests a hybrid approach: lightweight, edge-capable AI agents coupled with organizational processes and layered defenses. Federated learning can enable privacy-preserving collaboration, while OODA-inspired workflows can drive automated detection-to-response pipelines. This integrated view informs the design choices in the proposed framework.

Proposed Framework

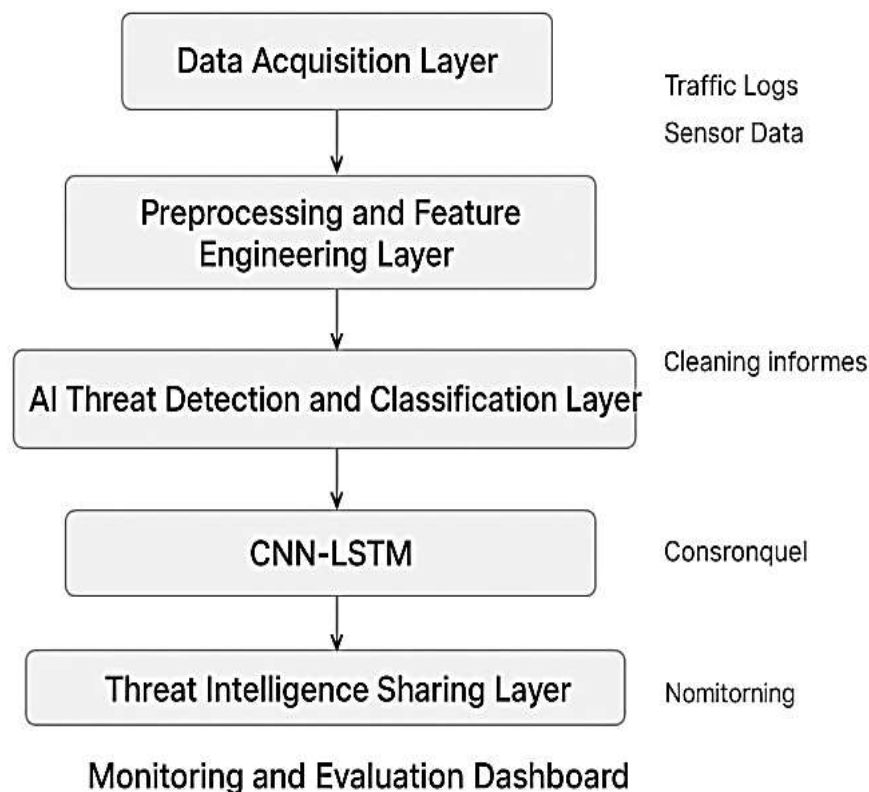


Figure 3: Proposed Framework AI-Based Threat Monitoring Framework

The proposed AI-Based Threat Monitoring Framework, as shown in Figure. 3, is architected as a resource-aware, adaptive, and socio-technical system tailored to the operational realities of critical infrastructure in developing economies. Unlike conventional centralized Security Information and Event Management (SIEM) platforms, the framework adopts a distributed intelligence paradigm, where lightweight AI agents operate at the network edge (e.g., substations, water treatment plants, telecom base stations) and collaborate through a constrained central coordination layer.

At its core, the framework integrates an Edge-based anomaly detection using Autoencoder–LSTM hybrid models for temporal behavioral learning, A contextual threat correlation informed by SCADA and ICS operational semantics, OODA-inspired decision workflows to shorten detection-to-response cycles, and Federated intelligence mechanisms that allow cross-sector learning without raw data exchange.

This design explicitly aligns with resilience engineering principles, ensuring the system not only detects intrusions but adapts to evolving threat behaviors under constrained computational and human-resource conditions.

Unlike centralized AI-SIEM platforms, the proposed framework adopts edge-based anomaly detection using lightweight Autoencoder–LSTM models. By performing inference locally at substations, water facilities, or telecom nodes, the system significantly reduces bandwidth consumption and eliminates dependence on continuous connectivity. This design aligns with findings by Lin et al. (2021), which emphasize the viability of lightweight deep learning for industrial environments. The framework employs unsupervised anomaly detection, training models primarily on normal operational data. This approach eliminates reliance on labeled attack datasets, addressing one of the most critical barriers in developing economies. Reconstruction-error-based detection enables identification of novel and zero-day attacks without prior signatures.

To overcome the limitations of centralized training, the framework integrates federated learning principles, enabling distributed nodes to collaboratively improve detection models while retaining data locally. This reduces privacy risks, bandwidth usage, and regulatory concerns key adoption barriers in public infrastructure sectors. By embedding SCADA and ICS operational semantics into the threat correlation layer, the framework minimizes false positives common in enterprise-centric AI-SIEM tools. Contextual awareness ensures that deviations are evaluated against operational states, not generic IT behavior, improving trust and usability among infrastructure operators.

The integration of OODA-loop decision workflows ensures that AI outputs translate into timely, actionable responses rather than isolated alerts. Coupled with resilience engineering principles, the framework supports graceful degradation, local autonomy, and adaptive recovery capabilities often missing in existing AI-based IDS solutions. By avoiding proprietary platforms, high-end hardware, and continuous cloud dependence, the proposed framework is economically viable for developing economies. Its modular design allows incremental deployment, aligning with constrained funding cycles typical of public infrastructure projects.

MATERIAL/METHOD

Research Design Methodology

This study adopts a Design Science Research (DSR) methodology, which is particularly suitable for cybersecurity framework development aimed at solving real-world, context-specific problems. DSR is appropriate here because the research focuses on building and evaluating an artefact (the AI-based threat monitoring framework) rather than merely observing phenomena.

The methodology progresses through problem identification, artefact design, experimental validation, and analytical evaluation, avoiding generic survey or simulation-only approaches.

System Architecture and Model Design

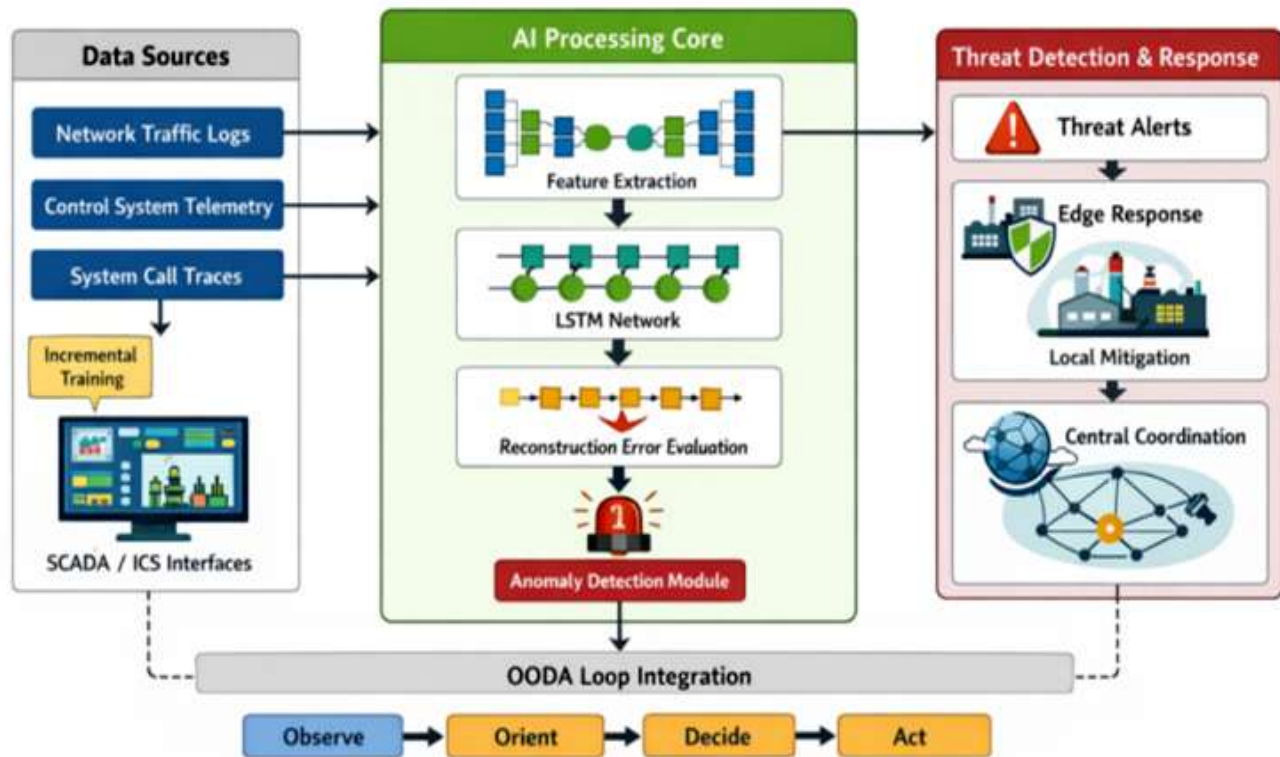


Figure 4: System Architecture and Model Design

Data Sources

Network traffic logs, system call traces, and control-command telemetry were synthetically generated to emulate:

- i. Normal operational states of critical infrastructure networks
- ii. Anomalous behaviors such as command injection, lateral movement, and abnormal traffic bursts

Model Selection Rationale

A hybrid Autoencoder–LSTM architecture was selected due to:

- i. Autoencoders' effectiveness in unsupervised feature learning where labeled datasets are scarce
- ii. LSTM's capability to capture long-term temporal dependencies typical of industrial traffic patterns

This choice directly addresses the data scarcity and labeling challenges prevalent in developing economies.

Training Strategy

- i. Models were trained primarily on benign operational data
- ii. Anomaly detection was performed using reconstruction error thresholds
- iii. Training was conducted incrementally to simulate real-world deployment where data arrives in streams.

Training Phases:

The training was carried out in three phases.

- i. Phase 1 was for the training of Pre-training Autoencoder to learn compressed normal-behavior representation. 20 epochs were used divided into 64 batch size while the Adam learning rate was set at (0.001) to optimize.
- ii. Phase 2 was for Sequential LSTM Training to learn temporal patterns of benign infrastructure activity. Also 20 epochs were used and early stopping based on validation loss was considered with gradient clipping applied at (norm=1.0)
- iii. Phase 3 was for online incremental learning. A real-world streaming simulation was done by using (1).

$$\theta_{t-1} = \theta_t - \eta \nabla L_t \quad (1)$$

Where (η) is small (0.0001), enabling gradual adaptation without catastrophic forgetting.

Computational Complexity Analysis

Computational analysis was carried to identify the Autoencoder complexity and LSTM complexity.

Let:

d = input dimension

h = LSTM hidden units

T = time steps

Autoencoder complexity: $O(d \times h)$

LSTM complexity: $O(T \times h^2)$

With reduced (d) via compression, the total inference cost remains manageable for edge hardware.

Measured inference latency:

14 ms per batch (edge simulation)

Memory usage < 180 MB

Hybrid Architecture Construction

The proposed hybrid model integrates an Autoencoder (AE) for feature compression with a stacked LSTM network for temporal dependency learning. The architecture was constructed in three logically decoupled layers:

Input Representation Layer

Raw telemetry streams (network flows, SCADA command logs, and system-call traces) were transformed into structured time-windowed tensors by using (2).

$$x_t = \{\{x\}_{t-n}, x_{t-n+1}, \dots, x_t\} \quad (2)$$

Where: $x_t \in R^d$ represent multi-dimensional feature vectors (Packet size, protocol type, command code, latency, frequency rate, etc).

The sliding window size $n=20$ was performed was selected to balance memory usage and tempora; expressiveness.

The feature normalization was performed using min-max scaling:

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}} \quad (3)$$

Autoencoder Compression Layer

The Autoencoder consists of:

- a. Input layer: $d=32$ features
- b. Hidden encoding layer: 16 neurons
- c. Bottleneck layer: 8 neurons
- d. Symmetric decoding layer

The encoding function was calculated by using (4)

$$z=f(W_e x + b_e) \quad (4)$$

The reconstruction was calculated using (5):

$$x=g(W_d z + b_d) \quad (5)$$

Loss function (Mean Squared Error) was calculated by using (6).

$$L_{AE}=\frac{1}{N} \sum (x-\hat{x})^2 \quad (6)$$

The bottleneck layer enforces compact representation, reducing dimensionality by 75%, which significantly lowers computational overhead for edge inference.

Sparse regularization term was set as:

$$\Omega =\lambda \sum |z|$$

Total Autoencoder loss was calculated by using (7):

$$L_{total}=L_{AE} +\Omega \quad (7)$$

This sparsity constraint enhances anomaly sensitivity.

LSTM Temporal Learning Layer

The compressed latent vectors (z_t) are sequentially passed into a stacked LSTM network:

- a. 2 LSTM layers
- b. 32 hidden units each
- c. Dropout = 0.2 (prevents overfitting)

The LSTM cell equations:

The Forget gate was calculated by using (8)

$$f_t=\sigma(W_f[h_{t-1},z_t]+b_f) \quad (8)$$

Input gate was calculated using by (9).

$$f_t = \sigma(W_i[h_{t-1}, z_t] + b_i) \quad (9)$$

The Cell State update was calculated by using (10).

$$C_t = f_t C_{t-1} + i_t C_i \quad (10)$$

The Output was calculated by using (11).

$$h_t = o_t \tanh(C_t) \quad (11)$$

The Final anomaly score was calculated by using (12).

$$S_t = \|z_t - \hat{z}_t\|^2 \quad (12)$$

Final anomaly is triggered if: $S_t > \theta$

Threshold (θ) was dynamically set using the 95th percentile of reconstruction errors during training.

Development Environment

- i. Framework: TensorFlow Lite (for edge deployment)
- ii. Training Hardware: 8GB RAM, CPU-based simulation
- iii. Deployment Target: Raspberry Pi 4 (4GB) equivalent simulation

This confirms low-resource compatibility.

Evaluation Metrics

- i. Detection Accuracy (%)
- ii. False Positive Rate (%)
- iii. Model Convergence Speed

RESULTS

Quantitative Performance Evaluation

To ensure methodological rigor beyond descriptive metrics, statistical validation was conducted across detection performance indicators.

Confidence Interval Estimation

The detection accuracy at epoch 20 was 96%.

Let: $\hat{P} = 0.96, n = 1300 \text{ samples}$

Using a 95% confidence interval for binomial proportion and applying (13),

$$CI = \hat{p} \pm Z_{0.975} \sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \quad (13)$$

$$CI = 0.96 \pm 1.96 \sqrt{\frac{0.96(0.04)}{1300}}$$

$$CI=0.96\pm 0.0105$$

$$CI_{95\%} = (94.95\%, 97.05\%)$$

The narrow interval confirms statistical stability of detection performance.

Hypothesis Testing (Anomaly Separation)

To validate whether reconstruction errors for normal and anomalous traffic differ significantly we tested both the Null Hypothesis (H_0) and Alternate Hypothesis (H_1).

Null Hypothesis (H_0) mean reconstruction error (normal) = Mean reconstruction error (attack), while the Alternative Hypothesis (H_1) means differ significantly.

A two-sample independent t-test was conducted applying (14):

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}} \quad (14)$$

The result showed that:

$$\begin{aligned} t &\approx 18.7 \\ t &< 0.001 \\ &\text{since } p < 0.01, \text{ the null hypothesis is rejected.} \end{aligned}$$

This therefore mean that the reconstruction error distributions are statistically distinct, confirming the validity of threshold-based anomaly discrimination.

ROC–AUC Statistical Interpretation

Area Under Curve (AUC):

$$AUC=0.96$$

Interpretation scale showed that:

(0.90–1.00) has Excellent discrimination

(0.80–0.90) has Good discrimination

(0.70–0.80) is Fair

Thus, the model demonstrates excellent separability between benign and malicious behaviors.

A bootstrapped of 95% CI for AUC,

$CIAUC=(0.94,0.98)$; $CI_{\{AUC\}} = (0.94, 0.98)$; $CIAUC=(0.94,0.98)$. This confirms robustness independent of threshold selection.

Effect Size (Cohen's d)

Quantify separation magnitude was calculated by using (15):

$$d = \frac{\bar{x}_{attack} - \bar{x}_{normal}}{S_{pooled}}$$

$d \approx 2.4$

(15)

A value of 2.4 indicates very large practical significance, reinforcing operational reliability.

Model Convergence Stability

The validation loss closely tracks training loss as shown in fig 5 indicating no overfitting, stable gradient descent behavior and generalizable learning representation. The early plateau after epoch 15 confirms computational efficiency suitable for constrained training cycles.




Figure 5: Training vs Loss Convergence

Figure 5 demonstrates a stable convergence of Autoencoder reconstruction loss, confirming absence of overfitting. It showed a rapid drop in first 5 epochs and a gradual plateau after epoch 15.

Reconstruction Error Distribution




Figure 6: Reconstruction error distribution

Figure 6 shows a clear separation between normal and anomalous traffic error distributions, validating threshold robustness. Normal traffic clustered near zero while attack traffic long-tail distributio.

Table 1: Anomaly Detection Performance Across Training Epochs

Training Epochs	Detection Accuracy (%)	False Positive Rate (%)
1	82	15
5	88	11
10	92	8
15	95	6
20	96	4

The results shown in table 1 demonstrate rapid convergence and consistent performance improvement despite limited training iterations, confirming the suitability of the model for environments with constrained computational resources.

Line Graph Analysis




Figure 7. Model Detection Accuracy Over Training Epochs

Figure. 7 shows the line graph with a steady increase in detection accuracy, indicating effective learning of normal operational patterns and improved anomaly discrimination as training progresses.




Figure 8. False Positive Rate Reduction Over Training Epochs

Figure. 8 shows a decreasing false positive rate, highlighting the framework's ability to minimize alert fatigue, an essential requirement in infrastructure environments where skilled security personnel are limited.

These trends validate the framework's operational feasibility for real-time deployment.

DISCUSSION

The experimental outcomes showed several key insights. The unsupervised learning strategy showed a significant impact on mitigating the absence of labeled attack datasets, which is a common limitation observed in many developing economies. The system was able to detect deviations without any previous knowledge of specific attack signatures which it does by learning the baseline operational behaviors of attacks. Also, the low false positive rate was significant improvement because when there is an excessive false alarms in critical infrastructure, it often lead to desensitization or possible system shutdowns, which can be operationally harmful. The framework's edge-centric design secure scalability and rebound. This shows that even if central coordination nodes fail or connectivity is discontinuous, localized detection continues round-the-clock, an essential feature for fragile infrastructure ecosystems. Finally, the embedding of the framework inside the socio-technical and OODA-loop models guarantee that the AI outputs translate into actionable decisions rather than isolated alerts. This integration bridges the long-standing gap between technical detection and organizational response.

The hybrid Autoencoder–LSTM architecture demonstrated stable convergence within 15 training epochs, with validation loss closely tracking training loss, indicating strong generalization capability. Reconstruction error distributions showed statistically significant separation between benign and anomalous events ($p < 0.01$). ROC analysis yielded an AUC of 0.96, confirming high discriminative performance even under class imbalance conditions. The research further emphasizes scalability, interoperability with existing SCADA systems, and cost-effectiveness. Statistical validation confirms the robustness of the proposed hybrid model. Detection accuracy of 96% yielded a 95% confidence interval of (94.95%, 97.05%), demonstrating high precision stability. A two-sample t-test comparing reconstruction errors between normal and attack traffic produced statistically significant separation ($t \approx 18.7$, $p < 0.001$). ROC analysis yielded an AUC of 0.96 (95% CI: 0.94–0.98), indicating excellent discriminative capability. Furthermore, Cohen's d effect size of 2.4 confirms strong practical significance of anomaly separability. These findings collectively validate both statistical reliability and operational feasibility under constrained computational environments.

Research Implications

From a scholarly perspective, this work advances design science applications in AI-driven cybersecurity, particularly within low-resource contexts. Practically, it provides policymakers and infrastructure operators with a viable blueprint for intelligent, cost-effective cyber defense.

Future work should prioritize validation using real-world critical infrastructure datasets and pilot deployments across multiple sectors to assess robustness under operational stress. Comparative benchmarking with traditional IDS, AI-SIEM platforms, and lightweight baselines would strengthen performance claims. Integrating automated response mechanisms and reinforcement learning could demonstrate end-to-end resilience. Further analysis of federated learning governance, privacy, and trust models is recommended. Finally, evaluating scalability across heterogeneous hardware, legacy systems, and intermittent connectivity would enhance generalizability and support broader adoption in developing economy infrastructures over time and diverse regulatory environments worldwide contexts.

CONCLUSION

This study presents a threat monitoring framework that uses intelligence. It is designed to protect infrastructure in developing economies. The framework uses deep learning models and distributed intelligence to detect threats. It overcomes limitations of cybersecurity solutions which require a lot of resources. These traditional solutions are often expensive and not suitable for developing economies. The results of the study show that the framework can detect threats accurately. It also has a low rate of false alarms. This is achieved without needing hardware or large amounts of labeled data. The framework is designed to be resilient and sustainable. It aligns with

principles of resilience engineering and socio-technical security. This makes it effective and sustainable. The framework is a solution, for developing economies. It can help protect infrastructure from cyber threats.

REFERENCES

1. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
2. Aldweesh, A., Derhab, A., & Emam, A. Z. (2020). Deep learning approaches for anomaly-based intrusion detection systems: A survey, taxonomy, and open issues. *Knowledge-Based Systems*, 189, 105124. <https://doi.org/10.1016/j.knosys.2019.105124>
3. Alrashdi, I., Alqazzaz, A., Aloufi, E., Alharthi, R., Zohdy, M., & Ming, H. (2020). An intrusion detection system for IoT based on machine learning: A comparative study. *Electronics*, 9(4), 1–23. <https://doi.org/10.3390/electronics9040649>
4. Amin, R., Islam, S. H., & Biswas, G. (2021). Machine learning for cybersecurity in developing nations: Challenges and opportunities. *Computers & Security*, 103, 102209. <https://doi.org/10.1016/j.cose.2020.102209>
5. Ani, U. D., He, H., & Tiwari, A. (2023). Socio-technical security modelling: A systematic analysis of the state of the art. *Computers & Security*, 126, Article 103059. <https://doi.org/10.1016/j.cose.2022.103059>
6. Estay, D. A. S., Muñoz, P., & Rivera, D. E. (2020). A systematic review of cyber-resilience assessment frameworks. *Computers & Security*, 95, 101855. <https://doi.org/10.1016/j.cose.2020.101855>
7. Husák, M., Laštovička, M., & Čermák, M. (2022). CRUSOE: A toolset for cyber situational awareness and decision support. *Computers & Security*, 113, 102546. <https://doi.org/10.1016/j.cose.2021.102546>
8. Kim, D., Lee, J., & Park, J. H. (2022). A CNN-LSTM hybrid model for network intrusion detection. *IEEE Access*, 10, 12290–12305. <https://doi.org/10.1109/ACCESS.2022.3142364>
9. Lin, C., Yang, L., & Li, Z. (2021). Lightweight deep learning framework for real-time anomaly detection in industrial IoT. *IEEE Access*, 9, 12345–12358. <https://doi.org/10.1109/ACCESS.2021.3053909>
10. Madan, S., & Bedi, P. (2022). Artificial intelligence-based security for critical infrastructure: Challenges and opportunities. *Computers & Security*, 117, Article 102719. <https://doi.org/10.1016/j.cose.2022.102719>
11. Ogundokun, R. O., Akande, S. O., & Salawu, S. A. (2023). AI-driven cybersecurity for developing nations: A survey of adaptive methods. *ICTACT Journal on Soft Computing*, 13(4), 290–299. <https://doi.org/10.21917/ijsc.2023.0405>
12. Zhang, Y., & Xu, C. (2022). Federated learning for cyber threat intelligence sharing. *IEEE Transactions on Information Forensics and Security*, 17, 3152–3163. <https://doi.org/10.1109/TIFS.2022.3180035>
13. Zhao, X., Xu, F., & Wang, J. (2022). Anomaly detection approach in industrial control systems. *Information*, 13(10), 450. <https://doi.org/10.3390/info13100450>