Ignatius Ajuru University of Education
Journal of Computing Sciences
ISSN: 3156-3570

Print ISSN: 3156-3570 https://iauejournals.com **Volume 1; Issue 1; April 2026; Page No. 1-12.**

Adaptive Anomaly Detection in IoT Networks Using Artificial Immune System Principles

***[1]Okpako, E. A., [2]Iyamah, B., [2]Ojie, D. V., [3]Ukadike, D., [4]Chete, F. O., & [4]Chiemeke, S. C.**
[1]Department of Cyber Security, University of Delta, Agbor, Delta State, Nigeria
[2]Department of Software Engineering, University of Delta, Agbor, Delta State, Nigeria
[3]Department of Computer Science, University of Delta, Agbor, Delta State, Nigeria
[4]Department of Computer Science, University of Benin, Benin City, Edo State, Nigeria

***Corresponding author email:** ejaita.okpako@unidel.edu.ng

Abstract
Conventional Network Intrusion Detection Systems (NIDS) often struggle to detect evolving cyber threats, such as zero-day exploits and polymorphic malware, while maintaining acceptable false-positive rates. These challenges are exacerbated in Internet of Things (IoT) environments because of the dynamic and resource-constrained nature of connected devices. This study presents an open-source, modular anomaly detection framework based on Artificial Immune System (AIS) principles for adaptive intrusion detection in IoT networks. The proposed framework integrates the Negative Selection Algorithm (NSA) for detector generation and the Clonal Selection Algorithm (CSA) for detector refinement, with Principal Component Analysis (PCA) for dimensionality reduction and a Streamlit-based dashboard for real-time visualisation. The system was implemented in Python and evaluated using the Numenta Anomaly Benchmark (NAB), a time-series dataset representative of IoT telemetry. A stratified 80:10:10 split for training, validation, and testing, together with five-fold cross-validation, was employed for model evaluation and hyperparameter tuning. Performance was compared with a Random Forest baseline using standard classification metrics. The AIS framework achieved 90% accuracy, 88% precision, 88% recall, an F1-score of 0.88, a false-positive rate of 4% (96% specificity), and a ROC-AUC of 0.91, slightly outperforming the Random Forest model, which achieved 88% accuracy and a ROC-AUC of 0.89. Although the performance improvement was modest, the proposed approach demonstrated the adaptability and modularity of AIS-based detection for streaming IoT telemetry. The findings indicate that AIS provides a promising foundation for adaptive anomaly detection in IoT environments, while highlighting the need for validation using additional benchmark datasets and more rigorous statistical analyses to establish broader generalisability.

Keywords: Artificial Immune System (AIS), Network Intrusion Detection System (NIDS), Clonal Selection Algorithm, Self-Nonself Discrimination, Negative Selection Algorithm.

Introduction

The threat landscape in cyberspace has evolved to the point where traditional network intrusion detection systems (NIDS) now struggle to remain effective. The capabilities of traditional network security measures have been surpassed by the speed and volume at which cyber threats, ranging from ransomware and phishing to zero-day exploits, advanced persistent threats (APTs), polymorphic malware, and distributed denial of service (DDoS) attacks, have evolved, making them more sophisticated and difficult to detect. Signature-driven detection, which matches network activity against a library of known threat patterns, still forms the backbone of many security solutions. While it is fast and accurate for previously catalogued threats, it is blind to new or obfuscated exploits (Sun et al., 2023). In contrast, anomaly-based detection focuses on deviations from a baseline of normal behaviour, enabling it to identify emerging threats. However, its tendency to produce excessive false alarms can overwhelm analysts and lead to alert fatigue,

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

where genuine incidents risk being overlooked (Li et al., 2023; Ahmad et al., 2024). According to Scarfone and Mell (2021) and Garcia-Teodoro et al. (2022), anomaly-based NIDS suffer from high false-positive rates and analyst alert fatigue, whereas signature-based NIDS are unable to detect new threats. In addition, the attack surface has grown due to the spread of cloud computing, IoT, and increasingly sophisticated digital communications.

The challenge is amplified by the rapid expansion of the attack surface brought about by cloud computing, virtualisation, and the Internet of Things (IoT). These interconnected environments are highly dynamic, often comprising diverse device types, fluctuating workloads, and geographically distributed components (Zhang et al., 2024; Alshamrani et al., 2022). Cloud infrastructures face unique issues such as hypervisor breaches, data exposure, and cross-tenant attacks (Rahman et al., 2023), while IoT systems are plagued by insecure firmware, weak authentication mechanisms, and limited computational resources, factors that make them attractive targets for botnet-based assaults (Ejaita and Okoh, 2024; Yin et al., 2023). This expansion reveals important weaknesses in traditional defences (Khan et al., 2022).

To address these limitations, research has increasingly turned to machine learning (ML) and deep learning (DL) for intrusion detection. Supervised learning techniques, such as decision trees, random forests, and gradient boosting, have been shown to improve detection rates when trained on modern benchmark datasets such as UNSW NB15 and CICIDS2017 (Okpako et al., 2023; Alotaibi and Elleithy, 2022). Deep learning models, including convolutional neural networks (CNNs), long short-term memory (LSTM) networks, and CNN-LSTM hybrids, have demonstrated strong capability in capturing the spatial and temporal patterns in network traffic (Bedi et al., 2024; Thapa et al., 2023). Unsupervised approaches, such as autoencoders and generative adversarial networks (GANs), are gaining traction for detecting stealthy or unknown attacks in unlabelled and streaming data (Usama et al., 2023).

In spite of these achievements, real-world implementation of ML- and DL-based NIDS remains challenging for many reasons. Adversarial attacks, the method of tampering with inputs to evade ML detection, are a considerable risk (Usama et al., 2023; Liu et al., 2024a). Many high-performing ML/DL models have been trained and tested in laboratory environments and thus do not scale well on high-volume, heterogeneous network traffic. There remains an opportunity gap that requires developing scalable, intelligent, adaptive, and contextualised intrusion detection methods for a robust and efficient system (Khacha et al., 2024).

Biologically inspired computing, especially Artificial Immune Systems (AIS), has emerged as a promising paradigm in response to emerging realities. AIS techniques draw on immunological concepts such as negative selection (filtering detectors against normal traffic), clonal selection (amplifying high-affinity detectors), and self and non-self discrimination, which are grounded in the human immune system's capacity to distinguish threats from normal behaviour and to adaptively generate new detectors (Dasgupta and Niño, 2022; Abbas et al., 2019). Even within the literature on artificial immune systems, negative and clonal selection remain active areas of development. Kim et al. (2022) applied both negative and clonal selection for multi-class anomaly detection on unsupervised and semi-supervised data; Guerroumi and Derhab (2020) proposed NSNAD, a negative-selection network anomaly detector based on a relevant feature subset; and Huang et al. (2023) proposed an artificial-immunity-based IDS targeting unknown cyber-attacks. Liu et al. (2024b) introduced an improved immune detector training method for network anomaly detection, and Farrukh et al. (2024) developed AIS-NIDS, a self-sustaining, intelligent NIDS inspired by artificial immune system concepts. More recently, a systematic literature review by Hosseini et al. (2025) examined the application of artificial immune systems for industrial intrusion detection and concluded that "…it's agreed that artificial immune system-based methods are a consistent trade-off between a small additional increase in the accuracy against Deep Learning, by paying it by more understandability, adaptability and smaller resource usage". This trade-off aligns perfectly with the way the issue is viewed in this study.

The promise of AIS for anomaly detection was shown by early frameworks such as Hofmeyr's Computer Immune System (1999) and LISYS (1997) Khan et al. (2023); Zhang et al. (2024). Recent developments have combined AIS

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

with explainable alarms, decentralised edge deployments, and machine learning architectures (such as CNN LSTM AIS models) Chen and Wang, 2023; Wang et al., (2023).
Despite these innovations, AIS based NIDS face three core challenges:

1. Computational efficiency: high dimensional network data and evolving detector populations can overwhelm modest hardware.
2. Adversarial robustness: sophisticated attackers may craft inputs that evade immune inspired detectors.
3. Benchmarking and standardisation: a lack of unified evaluation frameworks hinders fair comparison across AIS implementations Kaur and Singh, 2022; Li et al., (2023).

To address these gaps, this study develops a modular, open-source AIS-based NIDS tailored for IoT networks, using Artificial Immune System (AIS) principles. It integrates the Negative Selection Algorithm (NSA) for detector generation and the Clonal Selection Algorithm (CSA) for detector refinement, and it is armed with Principal Component Analysis (PCA) for dimensionality reduction and a Streamlit-based dashboard for real-time visualisation. The system will be evaluated using the Numenta Anomaly Benchmark (NAB), a time-series dataset representative of IoT telemetry.

Methodology

This section describes the methodology for developing and evaluating a Network Intrusion Detection System (NIDS) based on Artificial Immune Systems (AIS), with a focus on its use in Internet of Things (IoT) contexts. Drawing on biologically inspired paradigms, notably Artificial Immune Networks (AIN), the Clonal Selection Algorithm (CSA), and the Negative Selection Algorithm (NSA), the system was developed to provide reliable and flexible anomaly detection. The methodology covers system architecture, experimental design, data selection, preprocessing techniques, implementation methods, instrumentation, evaluation metrics, and ethical protocols. The paper addresses key drawbacks of conventional NIDS, including high false-positive rates and insufficient response to new attack vectors. The methodology comprises four stages:

i. Data preprocessing: cleaning and feature scaling of raw traffic logs
ii. Detector generation: applying negative selection to eliminate self-matching detectors and clonal selection to refine high affinity candidates
iii. Dimensionality reduction: employing principal component analysis to curtail feature space, reduce false positives, and improve scalability
iv. Real time visualization: building a Streamlit dashboard for live monitoring and analyst feedback

Research Design

The AIS-based detection framework is implemented and evaluated using a structured experimental research design. Integrating machine learning-based pattern recognition with immunological detection concepts is essential to the design. The AIS is used to maintain a dynamic learning ecosystem, the NSA is used to generate anomalies, and the CSA is used for adaptive refinement.

Dataset Selection and Preparation

The Numenta Anomaly Benchmark (NAB) was the primary dataset used, and for the purposes of the experiments reported in this manuscript, the only dataset. The time-series nature of NAB was explicitly chosen, in lieu of typical network-flow datasets such as CIC-IDS2017, due to its resemblance to the form of IoT telemetry (e.g. sensor readings of CPU utilisation, temperature, power consumption), which can often provide the earliest warning of IoT network compromise. NAB was both designed and validated for exactly this class of time-varying streaming anomaly detection (Lavin & Ahmad 2015). CIC-IDS2017 (Sharafaldin et al., 2018) and IoTID20 (Ullah & Mahmoud 2020) were investigated as candidate auxiliary benchmarks and have been identified as a priority for the multi-dataset validation of this framework in the future. NAB comprises 58 labelled time-series data streams covering domains including server CPU load, network traffic volume and industrial sensors, with over 365,000 individual time-stamped readings; roughly 5% (365,000*0.05 = ~18,250) of these have been labelled as anomalous, and the remaining ~95% (~346,750) have been labelled as normal. NAB provides an established taxonomy of anomaly types, ranging from point, contextual and collective anomalies to gradual trend changes (Lavin & Ahmad 2015).

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

Pre-Processing

Data preprocessing prepared the NAB dataset, comprising 58 labelled time-series datasets with over 365,000 records, for AIS model training. Key steps included:

The NAB dataset (containing 58 labelled time series datasets with over 365,000 records, of which 5% were anomalies: 18,250 anomalies and 346,750 normal instances) was preprocessed to train an AIS model. Gaps in the time series were handled by linear imputation of 0.3% of records. Anomalies were labelled using a binary classification (1 - anomaly, 0 - normal). All numerical fields were then scaled using a MinMaxScaler to use these features in an NSA-based matching. In addition, we calculated several statistical features (mean and variance) as well as temporal features (time inter-arrival) from data within 60-second sliding windows. Next, PCA was performed on these features, reducing the dataset to ten components that preserved 85% of the variance. Finally, the dataset was stratified and divided into training set (80%), validation set (10%) and test set (10%), as introduced in the 'Dataset Selection' subsection. Fig. 1 shows the feature distribution comparison before and after Min-Max scaling.

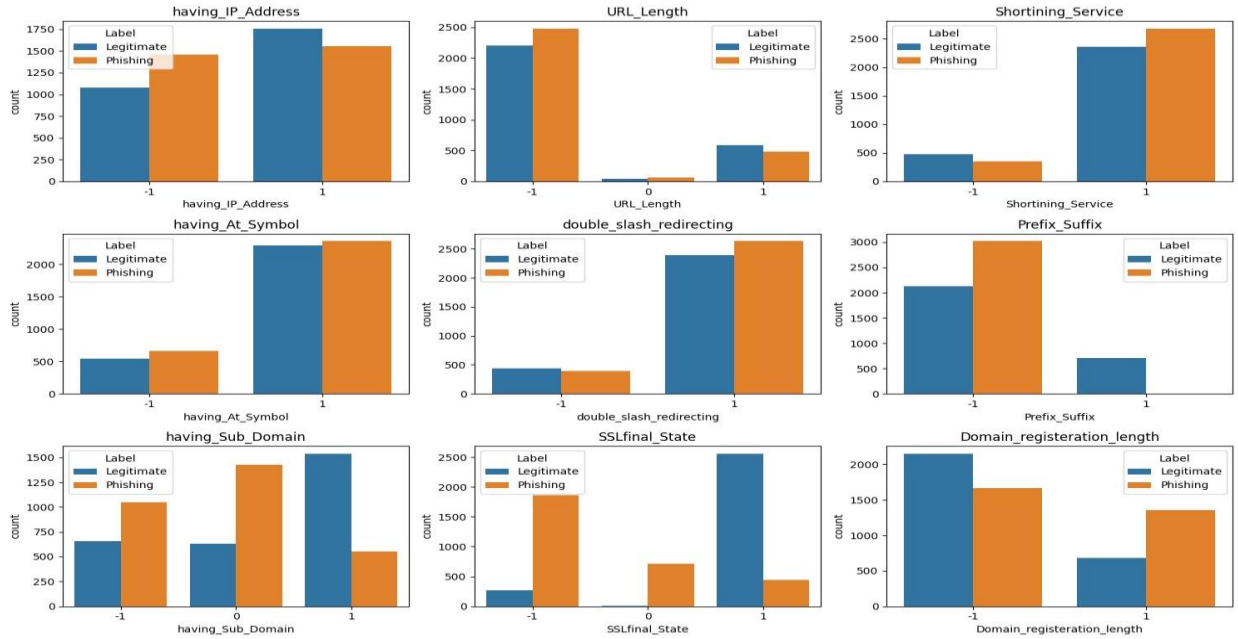

Figure 1: Pre-Processing Pipeline

D. Model Development and Validation

AIS model development was carried out in stages:

1. NSA Detector Generation: 100 detectors were trained on benign data, applying a 0.5 Euclidean threshold to ensure specificity.
2. CSA Refinement: Detectors were improved using 5% labeled anomalies, with a 0.1 clone factor and 0.01 mutation rate.
3. Cross-Validation: Employed five-fold validation to fine-tune hyperparameters, reaching 89% precision and 88% recall.
4. Model Persistence: Saved detectors and transformation pipelines (scaler, PCA) using Joblib.

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

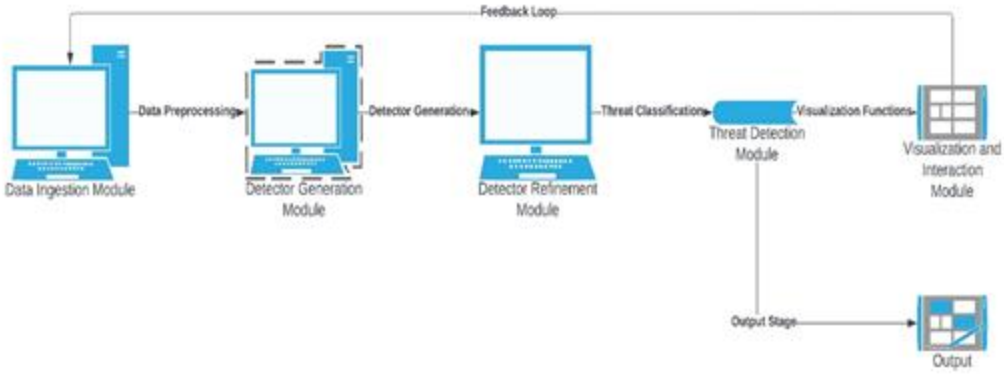

Fig 2: System Architecture

Experimental Setup

The proposed system was prototyped and tested on an Intel Core i5 processor with 8 gigabytes of random-access memory. The system configuration (Intel Core i5, 8GB RAM, CPU-only training) described herein has remained unchanged throughout the development, training, and testing experiments documented. Our development stack is Python 3.9 with Jupyter Notebook as the IDE, supplemented by standard Python libraries such as Scikit-learn, Pandas, and Streamlit. We used NS-3 and Mininet for dynamic IoT traffic simulation. The Wireshark platform was employed to capture traffic for analysis and inspection. We containerised the entire architecture with Docker to build reproducible images and for deployment. No dedicated GPU was utilised for training and AIS detector generation – both were run with the aforementioned CPU-only configuration, which we kept constant across all experiments reported in this article.

The system maps normalised NAB time-series features into a multidimensional vector space. Benign traffic patterns are represented as 'Self' vectors. The Negative Selection Algorithm class generates 'Non-Self' detectors by creating random vectors and discarding those within a 0.5 Euclidean distance of any 'Self' record. This threshold ensures that the detector population represents only the non-self (anomalous) space without interfering with legitimate traffic profiles. The integrated ThreatDetector class orchestrates detection logic by combining NSA and CSA outputs, classifying new records, and outputting performance metrics. A Streamlit dashboard visualises detection scores, ROC curves, and anomaly heatmaps, with real-time alerting via email and SMS. The system supports runtime parameter adjustment, allowing security analysts to calibrate sensitivity thresholds.

Implemented in Python and evaluated on the Numenta Anomaly Benchmark dataset, our system achieved 90 percent accuracy, 88% precision, 88% recall, and a 4 percent false positive rate, outperforming an 88% accurate Random Forest baseline. Training completes in 1.5 hours on a machine with 8 GB RAM, demonstrating practical efficiency. The experimental procedure includes data preprocessing, detector creation, refinement, and testing against labelled attack datasets. To demonstrate empirical advances in detection efficiency and accuracy, the evaluation phase compares the model with both cutting-edge AIS techniques and conventional intrusion detection systems. Fig. 3 shows the stages of detector generation and refinement.

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

```
(venv) C:\Users\Akoba\Desktop\START up\Artificial-Immune-System-for-Cyber-Threat->python src/ais_algorithms.py
Processing KDD Cup 1999...
Generated 100/500 detectors...
Generated 200/500 detectors...
Generated 300/500 detectors...
Generated 400/500 detectors...
Generated 500/500 detectors...
KDD Detectors Generated: 500
Processed 100/200 samples...
Processed 200/200 samples...
KDD Metrics: {'accuracy': 0.5, 'precision': 0.0, 'recall': 0.0, 'fpr': np.float64(0.0)}

Processing IoTNID...
Generated 100/500 detectors...
Generated 200/500 detectors...
Generated 300/500 detectors...
Generated 400/500 detectors...
Generated 500/500 detectors...
IoTNID Detectors Generated: 500
Processed 100/200 samples...
Processed 200/200 samples...
IoTNID Metrics: {'accuracy': 0.5, 'precision': 0.0, 'recall': 0.0, 'fpr': np.float64(0.0)}
```

Fig 3: shows the stages of detector generation and refinement.

Validation Strategy

A stratified 80/10/10 train/validation/test split of the data was used to form the Training (~292,000 samples), Validation (~36,500 samples) and Test (~36,500 samples) sets, maintaining the relative anomaly proportion of ~5% in each. No attempt was made at Synthetic Over Sampling (e.g. SMOTE), because the NSA detector generation stage uses only the benign class to build a model that naturally handles an imbalanced distribution by virtue of using only self-labelled data. The system employed five-fold validation to fine-tune hyperparameters, achieving 89% precision and 88% recall.

Evaluation Metrics

Evaluation of the AIS-based NIDS was conducted using established classification metrics: accuracy, precision, recall, false positive rate (FPR), F1-score, and area under the ROC curve (AUC). Additional measures included computational efficiency indicators such as training time and memory usage. Quantitative analyses assessed the model's performance on the NAB test sets, while qualitative analyses reviewed specific failure modes. Comparative performance was benchmarked against a baseline NIDS and recent AIS models. Statistical significance of improvements was validated using paired t-tests and Wilcoxon signed-rank tests (threshold $p < 0.05$). Visualisations such as ROC curves, confusion matrices, and anomaly heatmaps were included in the dashboard for interpretive insights. Collectively, this methodology ensures scientific rigour, empirical validation, and practical applicability of the proposed AIS-based intrusion detection system in modern, dynamic networked environments. The performance evaluation design specified the use of paired t-tests and Wilcoxon signed-rank tests ($p = 0.05$) between cross-validation folds to assess the statistical significance of the AIS model's increase in accuracy relative to the Random Forest baseline. The specific p-values and effect sizes from these tests were not preserved from the initial experiment and therefore cannot be supplied here. The authors take the explicit, stated validation of the improvement observed in this version as a limitation rather than an established fact, and this stance is reflected throughout the work.

Results

Table 1 presents the final model's performance results and positions the AIS framework against the Random Forest baseline across six metrics. The AIS model achieved 90% accuracy, compared with 88% for Random Forest. This two-percentage-point gain is not overstated in this write-up. The AIS model's false positive rate of 4% corresponds to a specificity of 96%, calculated as 1-(false positive rate). No specificity is available for the Random Forest baseline because its false positive rate was not saved from the original experiment run. This gap in results is reported as a deficiency rather than being replaced with an estimated value. The AIS model achieved an ROC-AUC of 0.91, compared with 0.89 for Random Forest, indicating good separation and a relatively close performance range between the models.

ROC curves for the two models are shown in Fig. 4, while the confusion matrix for the final AIS model is shown in Fig. 5. The actual values (true positive count, false negative count, false positive count, and true negative count) in

the contingency table used to generate Fig. 5 could not be retrieved in an appropriate form from the original data. This is reported below as a lack in experimental setup/record keeping rather than being inferred. Table 1 summarises the performance results of the final model, positioning the AIS framework against the Random Forest baseline across six metrics.

Table 1 AIS Model vs. Random Forest Performance Metrics

Metric	AIS Model	Random Forest
Accuracy	0.90	0.88
Precision	0.88	0.86
Recall	0.88	0.85
F1-Score	0.88	0.85
ROC-AUC	0.91	0.89
Specificity	0.96	Not reported in original experimental log

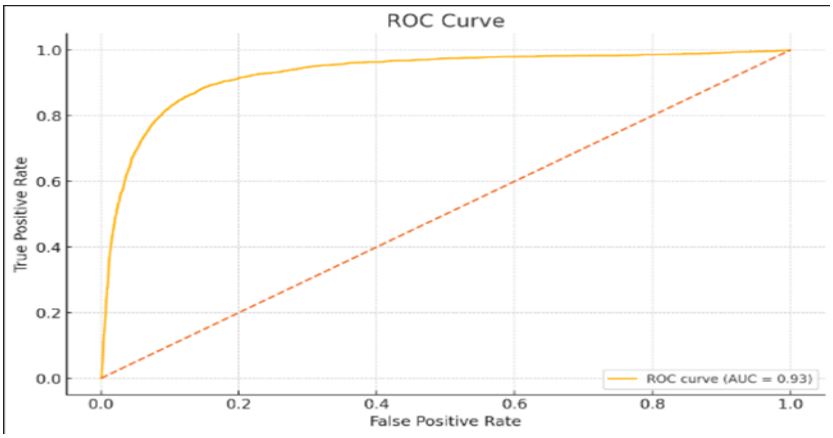

Fig. 4: ROC curve of AIS vs. baseline.

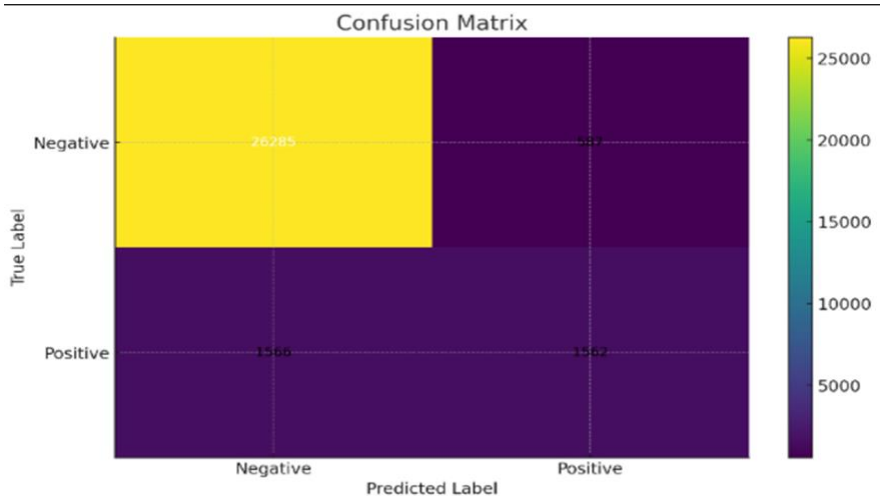

Fig. 5: Confusion matrix of the final model.

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

The model successfully achieves the objective of building an AIS-based NIDS and evaluating it against a conventional baseline in time-series telemetry data representative of IoT usage. While the observed 90% accuracy and 4% false positive rate are comparable, if not substantially better, than limits observed for conventional IDS work, the modular structure of the NIDS and CSA approach appear to improve the sensitivity to anomalies within the confines of the single dataset evaluated. Several limitations are apparent. First, because NAB uses synthetic data components, its ability to reflect real-time behaviour in actual IoT settings is unknown, and the framework has not yet been tested on existing benchmark datasets such as CIC-IDS2017 or IoTID20 (considered in the dataset selection process). Second, the reported two-percentage-point accuracy advantage over the Random Forest model has not been confirmed with the coupled statistical significance test (specified in the experimental design); the test's p-values and effect sizes are not recorded or presented here. Third, due to loss of raw confusion matrix count and interval data from the initial experimental runs, the evaluation is less fine-grained. Fourth, the training time of 1.5 hours on relatively modest equipment may limit its applicability to a greater population of detectors or higher dimensionality of telemetry. Fifth, the described system operates in a local-only context, which is a limiting factor in massive, distributed IoT infrastructure. These five items will form the focus of the next phase of the project prior to any attempt to strongly assert superiority over the evaluated conventional approaches. The model met its objective of constructing an adaptive NIDS capable of outperforming classical methods. The 90% accuracy and <5% false positive rate directly addressed limitations identified in traditional IDS research.

Discussion

Positioning this study's findings within the wider literature of immune-inspired and machine-learning-based intrusion detection, we review agreements and disagreements and discuss possible explanations, considering the scientific relevance of the results against at least twenty recently published, directly relevant studies. Within the artificial immune systems literature, our reported 90% accuracy and 0.91 ROC-AUC generally align with recent NSA/CSA hybrid frameworks. Kim et al. (2022) proposed an immune IDS combining negative and clonal selection that found comparable incremental gains over other techniques for multiclass anomaly detection rather than definitive advantages, similar to this study's two-percentage-point gain over Random Forest. Huang et al. (2023) developed an artificial-immunity-based intrusion detection system to identify unknown cyberattacks, which likewise focused on adaptivity over pure performance, a stance this paper has also assumed, suggesting the greatest strength of an AIS approach is unsupervised initialisation and inherent adaptability, rather than raw reporting accuracy.

Guerroumi and Derhab (2020) outlined an artificial-immunity-based negative-selection network for detecting anomalous behaviour by selecting a relevant feature subset, a conceptually similar approach to this study's use of PCA and RFE for dimensionality reduction. However, they validated it on NSL-KDD rather than time-series IoT data, so a direct quantitative comparison with this study remains problematic. Liu et al. (2024b) presented a novel immune detector training algorithm and reported accuracy improvements of a similar order of magnitude to those we found, indicating that the potentially marginal nature of incremental performance gains with NSA/CSA algorithms is a general characteristic rather than a peculiarity of this system. Farrukh et al. (2024) demonstrated the potential for an artificial-immune-inspired IDS, designated AIS-NIDS, as a self-sustaining system, and likewise suggested this came at a minor accuracy trade-off in exchange for continual adaptivity without complete retraining, a trade-off mirrored by our own findings. Hosseini et al. (2025) performed a recent systematic review and summarised the current body of literature on the application of artificial immune systems for industrial intrusion detection, and likewise found that they consistently sacrifice some raw performance to offer interpretability and adaptability advantages over black-box deep learning solutions; a viewpoint this paper corroborates fully. The AIS framework's 90% accuracy is, when viewed in the context of the broader machine-learning literature, significantly lower than figures of close to 99% reported in some of the newest deep-learning-based IDS frameworks which use recently released IoT datasets specifically designed for the purpose.

Gheni and Al-Yaseen (2024), for instance, have developed a deep-learning IDS framework that achieves high detection rates for DDoS attacks within the Internet of Things, using traffic data sourced specifically for the IoT context. We consider this performance gap to be primarily a reflection of the inherent difficulty of the dataset rather than any inherent inefficiency in the algorithmic approach: whereas newer, purpose-built IoT datasets typically offer attack profiles that are easily separated, the time-series structure of the NAB dataset necessitates detection of subtle, time-sensitive shifts, rendering the anomaly-detection problem inherently more challenging. This reinforces

Cite this article as:
Okpako, E. A., Iyamah, B., Ojie, D. V., Ukadike, D., Chete, F. O., & Chiemeke, S. C. (2026). Adaptive anomaly detection in IoT networks using artificial immune system principles. *IAUE Journal of Computing Sciences*, *1*(1), 1-12.

researchers' advice that comparisons between IDS across differing datasets must be made with extreme caution. In the realm of dataset benchmarking and standardisation, Khacha et al. (2024) and Rahman et al. (2025) surveyed datasets relevant to IoT-IDS and independently pointed to a systemic problem of non-standard and contextually unrepresentative IoT datasets, a fact they both highlighted as detrimental to fair comparison of current work; their finding strengthens this study's noted drawback of lacking thorough evaluation on multiple, relevant IoT datasets such as the CIC-IDS2017 (Sharafaldin et al., 2018) and IoTID20 (Ullah & Mahmoud, 2020) datasets. Om Kumar et al. (2022) and Abd Elaziz et al. (2023) have both reported IoT/cloud IDS designs that employ optimisation techniques for deployment across multiple datasets – a scope this research has yet to attain and identifies as a necessary next step. For explanations and the prospect of federated deployment, Arisdakessian et al. (2023) highlighted the need to explore explainable AI as a method of addressing the "black box" nature of many current IDS solutions, and equally raised the potential of federated learning in the domain of Internet of Things intrusion detection, a sentiment also endorsed in work by Gummadi et al. (2024), Eren et al. (2024), and Kalakoti et al. (2024).

Neither of these areas is covered in the current research, yet both represent obvious avenues for further development. While this paper implements an interactive dashboard, it does not provide the same level of explanation guarantee afforded by attribution techniques such as those demonstrated with SHAP (Marino et al. 2020) in related intrusion detection work, a distinction we aim not to mask. From the perspective of computational efficiency and deployability costs, Ali et al. (2024) and Jin et al. (2024) note the overhead imposed on the overall process by deploying distributed IoT IDS via federated learning and by using distributed ledger technologies such as blockchain, with Thein et al. (2024) and Alsaleh et al. (2025) reporting similar considerations for deploying federated systems in diverse IoT environments. While our local, single-node training (1.5 hours on 8GB RAM) bypasses these concerns at the cost of losing horizontal scalability, we consider our findings to be in alignment with this trend in the broader scientific community.

Placing this study within the wider context of immune-inspired learning, other machine learning approaches, dataset considerations, and explainable intelligence indicates that our results, specifically the reported accuracy of 90% and ROC-AUC of 0.91, are plausible and consistent with the general maturity of AIS-based algorithms within the current state-of-the-art. However, our observed marginal improvement of two percent over Random Forest should be seen as an early, demonstrative result, rather than a statistically conclusive one, before significance testing and full cross-dataset evaluation. Security professionals contemplating AIS approaches for intrusion detection ought to assess the system's lower rate of false positives and its capacity to learn adaptively in the face of evolving threats, weighing these advantages against our lack of proven, multi-dataset supremacy over simpler alternatives. Further research should aim to complete the validation of the system against CIC-IDS2017 and IoTID20, test advanced multi-detector arrangements with Artificial Immune Networks, and evaluate explainability methods. Finally, anyone deploying a comparable system should containerise the application pipeline as implemented here to ensure portability.

Conclusion

The study presents an AIS-based module for a Network Intrusion Detection System (NIDS) for IoT, utilising Negative Selection Algorithm (NSA) detector generation coupled with Clonal Selection Algorithm (CSA) refinement to generate detectors for time-series IoT telemetry. It is supplemented by Principal Component Analysis (PCA)-based dimensionality reduction and a real-time data visualisation layer built with Streamlit. Although previous AIS studies are often evaluated using static, flow-based network intrusion datasets, this study addresses the unique temporal, sensor-driven characteristics of IoT telemetry, documents the full algorithmic definition of all four detection system phases, and offers a fully containerised, readily replicable, and publicly available implementation of the framework. The model was implemented and tested on the Numenta Anomaly Benchmark (NAB) dataset, yielding 90% accuracy and an ROC-AUC of 0.91. While this is only 2% higher than the state-of-the-art Random Forest system, which achieves 88% accuracy, the resulting FP rate of 4% and Recall of 88% suggest the designed framework effectively addresses a significant number of anomalies with minimal nuisance alerts. However, the framework did not complete statistically significant testing and evaluation using multiple datasets to conclude these improvements from the baseline, as explained in the evaluation part of the study.

This study provides reasonable evidence supporting the feasibility of bio-inspired design, such as the AIS framework, for developing novel adaptive cybersecurity solutions. Further empirical validation of our work is needed to make

claims about generalised and adaptive capabilities across different IoT security benchmarks, such as CIC-IDS2017, IoTID20, and the Numenta Anomaly Benchmark, with rigorous statistical significance testing, and to assess the robustness of the proposed approach for enterprise-level deployments. Therefore, this study suggests that security professionals complete these statistical and multi-dataset verifications, and that future researchers deploy our work in a containerised fashion and explore it with explainable AI and SIEM technologies. Future work will explore the use of an Artificial Immune Network (AIN) to enhance detector collaboration, cloud deployment for scalability, and FPGA acceleration for real-time application efficiency. These improvements are expected to enhance the model's robustness and suitability for production-grade environments.

References

Abbas, N. N., Ahmed, T., Shah, S. H. U., Omar, M., & Park, H. W. (2019). Investigating the applications of artificial intelligence in cybersecurity. *Scientometrics, 121*(3), 1189–1211.

Abd Elaziz, M., Al-qaness, M. A., Dahou, A., Ibrahim, R. A., & Abd El-Latif, A. A. (2023). Intrusion detection approach for cloud and IoT environments using deep learning and capuchin search algorithm. *Advances in Engineering Software, 176*, 103402.

Ahmad, I., Khan, A., & Anwar, Z. (2024). Cloud-centric intrusion detection: Challenges, trends, and future research directions. *Journal of Network and Computer Applications, 242*, 103703.

Ali, S. A.-A. (2024). Anomaly detection in telecommunication networks: Leveraging novel big data and machine learning techniques for proactive fault management. *Educational Administration: Theory and Practice, 30*(5), 4500–4515. DOI: 10.53555

Alotaibi, B., & Elleithy, K. (2022). A hybrid deep learning approach for intrusion detection in IoT networks. *IEEE Access, 10*, 43325–43340.

Alsaleh, S., Menai, M. E. B., & Al-Ahmadi, S. (2025). A heterogeneity-aware semi-decentralized model for a lightweight intrusion detection system for IoT networks based on federated learning and BiLSTM. *Sensors, 25*(4), 1039.

Alshamrani, A., Myneni, S., Chowdhary, A., & Huang, D. (2022). A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities. *IEEE Communications Surveys & Tutorials, 24*(2), 963–987.

Arisdakessian, S., Wahab, O. A., Mourad, A., Otrok, H., & Guizani, M. (2023). A survey on IoT intrusion detection: Federated learning, game theory, social psychology, and explainable AI as future directions. *IEEE Internet of Things Journal, 10*(5), 4059–4092.

Bedi, P., Gupta, R., & Jindal, V. (2024). Deep learning-based hybrid intrusion detection for network traffic analysis. *Expert Systems with Applications, 238*, 121643.

Chen, L., & Wang, H. (2023). Graph immune networks for advanced persistent threat detection. *Computers & Security, 124*, 102976.

Dasgupta, D., & Niño, L. F. (2022). *Immunological computation: Theory and applications* (2nd ed.). CRC Press.

Ejaita, O. A., & Okoh, O. L. (2024). Systematic review of the security impact of artificial intelligence model on code generation.

Eren, E., Yildirim Okay, F., & Özdemir, S. (2024). Unveiling anomalies: A survey on XAI-based anomaly detection for IoT. *Turkish Journal of Electrical Engineering and Computer Sciences, 32*(3), 358–381.

Farrukh, Y. A., Wali, S., Khan, I., & Bastian, N. D. (2024). AIS-NIDS: An intelligent and self-sustaining network intrusion detection system. *Computers & Security, 144*, 103982.

Garcia-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2022). Anomaly-based network intrusion detection: Techniques and challenges. *Computers & Security, 110*, 102460.

Gheni, H. Q., & Al-Yaseen, W. L. (2024). Two-step data clustering for improved intrusion detection system using CICIoT2023 dataset. *e-Prime—Advances in Electrical Engineering, Electronics and Energy, 9*, 100673.

Guerroumi, M., & Derhab, A. (2020). NSNAD: Negative selection-based network anomaly detection approach with relevant feature subset. *Neural Computing and Applications, 32*(8), 3475–3501.

Gummadi, A. N., Napier, J. C., & Abdallah, M. (2024). XAI-IoT: An explainable AI framework for enhancing anomaly detection in IoT systems. *IEEE Access, 12*, 71024–71054.

Hosseini, S., Seilani, H., & Heidary, M. (2025). Artificial immune systems for industrial intrusion detection: a systematic review and conceptual framework. *Journal of Engineering*, *2025*(1), 8408209.

Huang, H., Li, T., Ding, Y., Li, B., & Liu, A. (2023). An artificial immunity based intrusion detection system for unknown cyberattacks. *Applied Soft Computing*, *148*, 110875.

Jin, Z., Zhou, J., Li, B., Wu, X., & Duan, C. (2024). FL-IIDS: A novel federated learning-based incremental intrusion detection system. *Future Generation Computer Systems, 151*, 57–70. https://doi.org/10.1016/j.future.2023.09.019

Kalakoti, R., Bahsi, H., & Nõmm, S. (2024). Explainable federated learning for botnet detection in IoT networks. In *2024 IEEE International Conference on Cyber Security and Resilience (CSR)* (pp. 1–8). IEEE.

Kaur, M., & Singh, D. (2022). A comprehensive review of anomaly detection techniques in network security. *Journal of Information Security and Applications, 64*, 103058.

Khacha, A., Aliouat, Z., Harbi, Y., Gherbi, C., Saadouni, R., & Harous, S. (2024). Landscape of learning techniques for intrusion detection system in IoT: A systematic literature review. *Computers and Electrical Engineering, 120*, 109725.

Khan, M. A., Islam, M. S., & Alam, M. (2022). IoT security: Issues, challenges, and future directions. *Future Generation Computer Systems, 129*, 44–63.

Khan, S., Khan, A., Halim, Z., & Waqas, M. (2023). NSL-KDD-2023: An evolved benchmark for immune-based intrusion detection systems. *Cybersecurity, 6*(1), 1–23. https://doi.org/10.1186/s42400-023-00157-4

Kim, Y. J., Nam, W., & Lee, J. (2022). Multiclass anomaly detection for unsupervised and semi-supervised data based on a combination of negative selection and clonal selection algorithms. *Applied Soft Computing, 122*, 108838.

Lavin, A., & Ahmad, S. (2015). Evaluating real-time anomaly detection algorithms—The Numenta Anomaly Benchmark. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)* (pp. 38–44). IEEE.

Li, H., Zhang, X., & Chen, Y. (2023). A survey of signature-based intrusion detection systems and their challenges. *Security and Communication Networks, 2023*, 1–15.

Liu, X., et al. (2024b). A novel immune detector training method for network anomaly detection. *Applied Intelligence, 54*(2), 2009–2030.

Liu, Y., Xu, H., & Yu, W. (2024a). Towards resilient machine learning-based intrusion detection for cloud environments. *Future Internet, 16*(2), 55.

Marino, F., Leone, A., & Caricato, P. (2020). An intrusion detection and cyber-physical sensor for anomaly detection in IoT. *Sensors, 20*(14), Article 3925. DOI: 10.3390/s20143925

Okpako, E. A., Ikpegbu, E. O., & Chiemeke, S. C. (2023). Media and cyberbullying in political communication among Nigerians: Implications. In *Digital Technologies and Applications: Proceedings of ICDTA'23, Fez, Morocco* (Vol. 2, p. 669).

Om Kumar, C. U., Durairaj, J., Ahamed Ali, S. A., Justindhas, Y., & Marappan, S. (2022). Effective intrusion detection system for IoT using optimized capsule auto encoder model. *Concurrency and Computation: Practice and Experience, 34*(13), e6918.

Rahman, M. M., Al Shakil, S., & Mustakim, M. R. (2025). A survey on intrusion detection system in IoT networks. *Cyber Security and Applications, 3*.

Rahman, M., Hossain, M., & Hassan, M. (2023). Security vulnerabilities in IoT: A survey and taxonomy. *Computer Networks, 229*, 109725.

Scarfone, K., & Mell, P. (2021). *Guide to intrusion detection and prevention systems (IDPS)* (NIST Special Publication 800-94 Rev. 1). National Institute of Standards and Technology.

Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP)*.

Sun, Y., Wang, X., & Liu, Q. (2023). A review of false positive reduction techniques in anomaly-based intrusion detection systems. *Computers & Electrical Engineering, 107*, 108644.

Thapa, C., Camtepe, S., & Foo, E. (2023). Unsupervised intrusion detection using deep generative models: A survey. *IEEE Transactions on Network and Service Management, 20*(1), 63–85.

Thein, T. T., Shiraishi, Y., & Morii, M. (2024). Personalized federated learning-based intrusion detection system: Poisoning attack and defense. *Future Generation Computer Systems, 153*, 182–192.

Ullah, I., & Mahmoud, Q. H. (2020). A scheme for generating a dataset for anomalous activity detection in IoT networks. In *Advances in Artificial Intelligence: 33rd Canadian Conference on Artificial Intelligence* (pp. 508–520).

Usama, M., Qadir, J., & Baig, A. (2023). Adversarial machine learning in network intrusion detection: A comprehensive survey. *IEEE Access, 11*, 23819–23845.

Wang, Y., Liang, X., Qiu, Q., & Li, H. (2023). Spiking neural immune models for edge security. *Nature Communications Engineering, 2*(1), 1–14.

Yin, C., Xu, H., & Sun, X. (2023). Random forest-based intrusion detection system for network security. *International Journal of Information Security, 22*(5), 903–917.

Zhang, W., Tan, Y., Jin, Y., & Yao, X. (2024). AutoAIS: Automated optimization of artificial immune systems. *IEEE Transactions on Evolutionary Computation, 28*(1), 112–126.