Development of a Hybrid Anomaly-Based Intrusion Detection System Using Autoencoder and Isolation Forest

Somtochukwu Chinedu (chinedusommie@gmail.com)[1], Andrew A Uduimoh (a.uduimoh@futminna.edu.ng)[1], Anyaora Peter (p.anyaora@futminna.edu.ng)[1], John k Alhassan (jkalhassan@futminna.edu.ng)[1], Hadiza Yusuf (hyusuf@umich.edu)[2]

[1]Cyber Security Department, Federal University of Technology Minna

[2]Computer and Information Science, University of Michigan-Dearborn

ABSTRACT

As cyberattacks advance in sophistication and fluidity, modern intrusion detection systems (IDS) must progress from static, signature-based models to adaptive models that can detect known and zero-day threats (Labonne, 2020; Ali et al., 2023). This study proposes a hybrid anomaly-based IDS that integrates an Autoencoder (AE) for deep feature representation with an Isolation Forest (IF) for statistical anomaly scoring. The objective is to enhance detection performance for both known and zero-day attacks within an unsupervised learning framework. The proposed model is an adaptation of the hybrid AE–IF framework devised by Mohammed and Telek (2023) and the fog-computing adaptation by Sadaf and Sultana (2020). The hybrid model deploys a logical fusion framework—Hybrid OR and Hybrid AND—dynamically balancing precision and recall in anomaly detection. The model is implemented in Python and trained using unsupervised learning on two benchmark datasets, NSL-KDD and CICIDS2017, so the model operates without any prior knowledge of attack signatures (Engelen et al., 2021). Experimental responses to model anomaly detection capabilities support that the Hybrid OR configuration of the model provided the most balanced anomaly detection performance scores recording F1-score of 0.85 and 0.69 on NSL-KDD and CICIDS2017 datasets respectively—both hybrid performance metrics outperforming the stand-alone AE and IF models. These results are in line with more recent evidence supporting the position that combining feature-reconstruction learning with statistical isolation increases the resilience and adaptability towards network intrusion detection (Elsaid & Binbusayyis, 2024;

Alhassan et al., 2024). While performance declined on the more complex CICIDS2017 dataset, the hybrid approach demonstrated improved generalisation relative to individual models. The results suggest that logical fusion of representation learning and isolation-based scoring provides a lightweight and adaptable framework for network intrusion detection, although further validation in live network environments is required before operational deployment (Abuabed et al., 2023; Wang et al., 2023).

Keywords: Anomaly Detection, Autoencoder, Cybersecurity, Hybrid Model, Intrusion Detection System, Isolation Forest

1. INTRODUCTION

The sophistication of cyberattacks has uncovered limitations with reactive, signature-based intrusion detection systems (IDSs)—their strength is detecting known patterns as threats, but unable to detect novel or multi-stage attacks (Labonne, 2020., Mirsky et al., 2018; Engelen et al., 2021). This limitation has motivated the growing adoption of anomaly-based detection approaches, which model normal network behaviour and flag deviations as potential intrusions. Recent studies have demonstrated that unsupervised learning techniques are particularly suitable for anomaly detection in environments where labelled data are scarce or rapidly outdated (Ali et al., 2023; Mohammed & Telek, 2023). In contrast to supervised classifiers that depend on predefined attack signatures, unsupervised models learn latent representations of normal traffic patterns and can therefore generalise to previously unseen threats. This property is especially important in dynamically evolving digital ecosystems, including financial and API-driven infrastructures, where attack vectors change frequently (Wang et al., 2023; Abuabed et al., 2023). Among unsupervised approaches, Autoencoders (AEs) and Isolation Forests (IFs) have gained significant attention. AEs provide compact, high-dimensional feature representations capable of modelling nonlinear network behaviour (Kimanzi et al., 2024; Adebayo et al., 2020). Isolation Forest, on the other hand, identifies

anomalies through random partitioning and exhibits linear scalability, making it computationally efficient for resource-constrained environments (Elsaid & Binbusayyis, 2024; Engelen et al., 2021). Despite these advantages, each model exhibits limitations when deployed independently. Isolation Forest may fail to capture subtle or highly correlated anomalies in complex feature spaces, while Autoencoders may struggle to detect sparse or statistically irregular events. Hybrid architectures have therefore been proposed to combine complementary strengths. In particular, AE→IF pipelines leverage representation learning for dimensionality reduction and apply statistical isolation for anomaly scoring, resulting in improved robustness against zero-day attacks and class imbalance (Intrusion Detection Based on AE and IF in Fog Computing, 2022; HIF, 2023). Prior studies have demonstrated performance improvements on benchmark datasets using either sequential or parallel hybrid configurations (Seq2Seq & ConvLSTM Hybrid, 2023; LSTM-AE Model, 2023). However, many of these implementations focus on single datasets or fixed fusion mechanisms, limiting insight into how logical decision strategies influence precision–recall trade-offs across heterogeneous traffic environments. This study addresses that gap by designing and evaluating a hybrid anomaly-based IDS that integrates Autoencoder-based feature extraction with Isolation Forest anomaly scoring and explicitly investigates two logical fusion strategies: Hybrid OR and Hybrid AND. The model is evaluated on two complementary benchmark datasets, NSL-KDD and CICIDS2017, enabling comparison between controlled experimental conditions and more complex, contemporary traffic distributions. By analysing performance trade-offs across both datasets, the study aims to provide clearer empirical evidence on how logical fusion influences detection sensitivity and false alarm rates within unsupervised IDS frameworks.

Objectives and Contributions. The study seeks to design and implement a sequential AE→IF hybrid intrusion detection system using unsupervised learning to detect both known and unknown attacks. Specifically, the study: (i) prepares and preprocesses the NSL-KDD and

CICIDS2017 datasets; (ii) constructs a hybrid Autoencoder–Isolation Forest architecture; (iii) evaluates performance using standard classification metrics including Accuracy, Precision, Recall, and F1-score; and (iv) analyses detection trade-offs under Hybrid OR and Hybrid AND decision rules. The principal contributions include a reproducible hybrid architecture, cross-dataset evaluation of logical fusion strategies, and empirical insights into precision–recall balancing within anomaly-based IDS design.

2. LITERATURE REVIEW

Over time, intrusion detection systems (IDSs) have transitioned from deterministic, rule based approaches, towards adaptive, hybrid architectures that combine statistical and neural learning approaches. Early IDS architectures relied on fixed signatures and statistical heuristics, which provided interpretability and low computational overhead but demonstrated limited capacity to detect zero-day or polymorphic attacks (Prabu & Sudhakar, 2023; Labonne, 2020). Although traditional machine learning classifiers such as Decision Trees, Random Forests, and Naïve Bayes improved detection precision under controlled datasets, their performance degraded when confronted with heterogeneous and evolving traffic patterns (Ali et al., 2023; Onyekpeze et al., 2021). Current research can be combined broadly into three areas: (1) rule-based and traditional machine-learning IDSs, (2) deep learning-based anomaly, and (3) hybrid Autoencoder-Isolation Forest models in which an Autoencoder learns a data representation, then an Isolation Forest classifier identifies anomalies. Together, these areas establish the knowledge base, and provide the rationale for the hybrid model presented in this dissertation.

2.1 Rule-Based and Traditional Machine Learning Approaches.

Comparative studies reinforce the limitations of classical methods. Ajagbe et al. (2024) demonstrated that ensemble learners such as XGBoost outperform single classifiers on UNSW-NB15, highlighting the importance of model aggregation for improved recall and precision. Similarly, unsupervised techniques including Isolation Forest and One-Class SVM have shown interpretability advantages, yet struggled with high-dimensional and imbalanced traffic distributions (IJIREEICE, 2022). Clustering-based approaches such as K-means and DBSCAN variants have also been associated with persistent false positives in dynamic environments (Fuhnwi et al., 2023; Kulkarni et al., 2022). These findings collectively suggest that while classical and algorithmic methods provide strong baselines, they lack robustness when deployed in modern, heterogeneous network ecosystems. Feature selection and optimisation techniques have been introduced to mitigate these weaknesses. For example, Alhassan et al. (2024) applied correlated-based feature selection to improve dimensional reduction, while Elsaid and Binbusayyis (2024) enhanced Isolation Forest performance using metaheuristic optimisation in industrial IoT contexts.

Although such improvements increase accuracy, they do not fully resolve generalisation limitations under evolving traffic conditions.

2.2 Deep Learning–Based Anomaly Detection.

The introduction of deep learning has significantly reshaped anomaly-based IDS design. Autoencoders (AEs), convolutional architectures, and recurrent models have demonstrated strong capacity for learning nonlinear feature representations (Kimanzi et al., 2024). Empirical evaluations using stacked and denoising AEs on NSL-KDD and CICIDS2017 reported high reconstruction fidelity and improved anomaly detection (Moraboena et al., 2020; Alrayes et al., 2024). More advanced temporal models such as LSTM-AE and Seq2Seq–ConvLSTM frameworks have further reduced false alarm rates by capturing sequential dependencies in traffic flows (Hnamte et al., 2023; Hariharan et al., 2025). However, these architectures often introduce increased computational complexity, which may limit deployment in resource-constrained or edge environments (Ali et al., 2023; Alsaleh et al., 2024). Isolation Forest (IF) remains attractive due to its linear scalability and efficiency in isolating outliers through random partitioning (Elsaid & Binbusayyis, 2024; Engelen et al., 2021). Nevertheless, IF alone may struggle to detect subtle anomalies embedded within high-dimensional latent spaces. This complementary relationship between representation learning and statistical isolation motivates hybridisation.

2.3 Hybrid Autoencoder–Isolation Forest Framework

Hybrid Autoencoder–Isolation Forest frameworks have demonstrated promising empirical improvements. Mohammed and Telek (2023) published an AE→IF pipeline that leads to improved accuracy against the ADBench datasets while demonstrating how AE latent embeddings can also improve IF separation. Sadaf and Sultana (2020) offered further expansion of this hybrid framework presented in fog computing with improved accuracy and inference speed under limited conditions. The Hybrid Isolation Forest (HIF) model extended this concept through distance-based calibration (Marteau et al., 2017; Kulkarni et al., 2022), and optimisation-based enhancements such as OIFIDS incorporated metaheuristic feature selection (Elsaid & Binbusayyis, 2024).

There is also hybrid interest in edge and federated contexts. Alhassan et al. (2024) made an application of AE feature compression to obtain 94.32% accuracy on NSL-KDD, while Alsaleh et al. (2024) proposed federated IDS architectures suitable for heterogeneous IoT. Kitsune (Mirsky

et al., 2018) demonstrated that an ensemble based on micro-autoencoders could perform online anomaly detection on Raspberry Pi class hardware supporting hybridisation can also be low-weight.

Threat modelling frameworks based on STRIDE (Abuabed et al., 2023; Prabu & Sudhakar, 2023) have also been integrated with designs of ML IDS to enhance interpretability and risk alignment. Das et al. (2024) additionally extended STRIDE by incorporating DREAD and SAHARA, in the context of automotive HPC assessment, to show that contextual threat modelling can inform ML-based intrusion defences. Taken together, this body of work comes to a shared conclusion: AE-IF hybrids outperform single-model baselines without compromising scalability to IoT, fog, and federated deployments.

2.4 Research Gap.

Although past studies have established that hybrid models can outperform standalone models, there are still some restrictions to the research done on these models. First, research into cross-dataset generalisation is usually not done thoroughly. Secondly, it is common to assume a logical fusion method without a systematic analysis of the logical fusion method. Lastly, discussions comparing the performance of different models typically only compare accuracy and do not discuss the trade-off of an actual false alarm with the performance of the model. Therefore, this study aims to fill these gaps by evaluating a sequential architecture using an AE to IF hybrid model and analysing the performance of that architecture across the NSL-KDD and CICIDS2017 datasets, using Hybrid OR and Hybrid AND as the logical fusion methods, respectively. The study contributes more precise empirical data regarding the performance of hybrid anomaly detection systems through examining precision-recall balancing based upon both controlled and complex traffic conditions.

3. METHODOLOGY

The methodology was based on a well-structured and easily reproducible process with three major steps: Dataset Description, Model Architecture, and Hybrid Strategy with Evaluation Metrics. The process was designed to be rigorous and easily reproducible with respect to benchmark data sets.

3.1 Dataset Description

The proposed hybrid intrusion detection method was trained and tested on two benchmarks intrusion detection data sets, namely NSL-KDD and CICIDS2017.

NSL-KDD was obtained from the original KDD ’99 data set, after eliminating redundancy and imbalance in class data. The data set consists of 125,973 training examples and 22,544 testing examples, and each of them has 41 features, including basic, content-based, and traffic features. The data set is divided in terms of attacks: Denial of Service (DoS), Probe, Remote-to-Local (R2L), and User-to-Root (U2R), giving it a well-structured environment for controlled experimentation

CICIDS2017 was downloaded from the Canadian Institute of Cybersecurity and was characterized by contemporary enterprise traffic with over 2.8 million flows and 80 features.

Three representative subsets were chosen based on computational convenience, including Monday (benign), Wednesday (web attacks), and Friday (DDoS attacks).

Preprocessing included data cleaning, label encoding, one-hot encoding for categorical features, and MinMax scaling of numerical attributes. MinMax scaling was selected to normalise feature ranges into a bounded interval [0,1], which improves neural network convergence stability and prevents dominant feature magnitude effects during Autoencoder training. For CICIDS2017, a stratified 70:30 train–test split was applied to preserve class distribution across partitions.

The dual-dataset design enables comparative evaluation between a controlled legacy dataset (NSL-KDD) and a higher-dimensional contemporary dataset (CICIDS2017), thereby facilitating assessment of cross-dataset generalisation.

3.2 Model Architecture

The proposed model incorporated two unsupervised machine learning algorithms—Autoencoder (AE) and Isolation Forest (IF)—into a sequential model combination paradigm.

Autoencoder

A symmetrical feedforward neural network that includes three encoder layers (32, 16, 8 neurons) and a mirrored decoder was selected for implementation as the Autoencoder. The reduced dimensions of 41 and 80 dimensions of input feature space into a compressed latent representation of dimension 8 has been designed to produce an 8-dimensional representation for

continual dimensional reduction. The size of 8 was chosen to create an appropriate balance with sufficient representation capability and lower dimensionality, ensuring compact embeddings that still preserved information about the underlying relationship of the features. A latent space that is too small risks losing a piece of information, while a latent space that is too large would produce less compression advantages, and potentially reduce the classifiers ability to separate the anomalies.

For encoder layer activation functions, a ReLU was selected to produce sparse representations in the case of vanishing gradient issues, during training. For all normalised input feature reconstruction, the decoder will use the Sigmoid activation function, since it will ensure that the reconstructed features maintain their scaling within [0,1] using MinMax scaling. The Autoencoder was solely trained with normal instances of traffic behaviour to learn how the baseline behaviour of normal traffic should display. The reconstruction loss function used to perform the reconstruction of input features with continuous values was MSE (Mean Squared Error) as MSE was very sensitive to deviation magnitude & was the best metric, given than the purpose of the reconstruction loss function was to reconstruct an input value with continuous characteristics.

The reconstruction error will be used to determine if given instance of traffic is an anomaly by calculating the reconstruction errors for the training data (normal) behaviour, producing a distribution of reconstruction errors from which the threshold will be a percentile based assessment. Specifically, the threshold was selected at the upper percentile of the normal reconstruction error distribution to limit false positives while maintaining sensitivity to deviations. This approach avoids arbitrary thresholding and aligns with established anomaly detection practices in unsupervised settings. The structure of the Autoencoder used for feature reconstruction is illustrated in Figure 1.

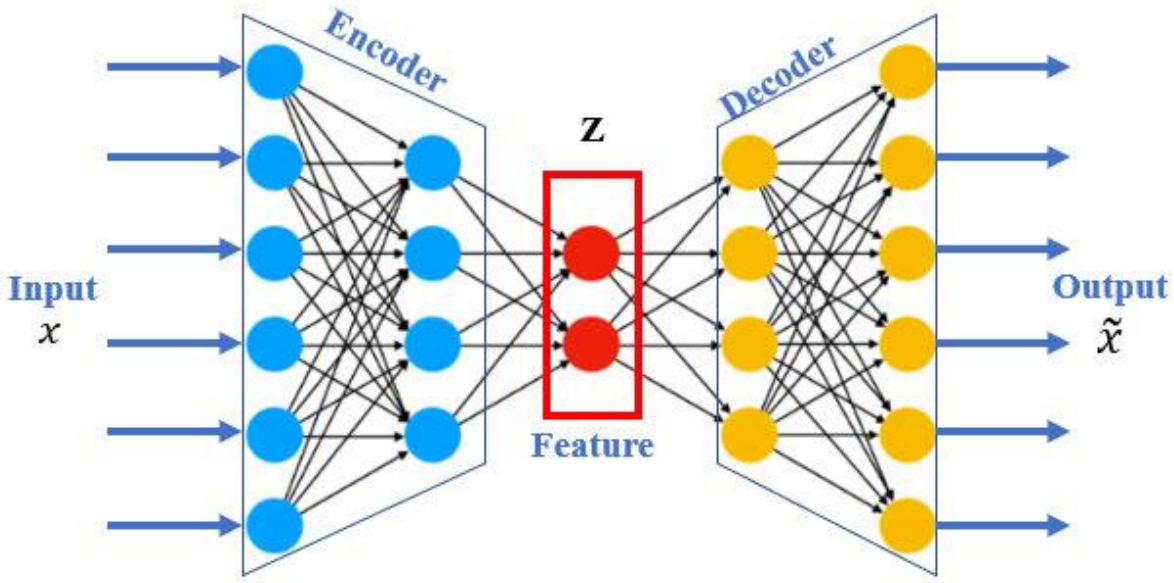

Figure 1: Encoder and Decoder Architecture of the Autoencoder Model (Alrayes et al., 2024)

The Autoencoder reduces the high-dimensional network input (x) to a latent space representation (z) via the encoder and decompresses it to reproduce the input as the output ($\bar{x}$) via the decoder. The difference between the network input and the reconstructed output $\bar{x}$ is the anomaly score to identify abnormal network traffic patterns.

Isolation Forest.

The Isolation Forest component was trained on the 8-dimensional latent representations produced by the encoder. Operating in the compressed latent space reduces dimensional complexity and enhances isolation sensitivity by concentrating discriminative features. The number of estimators was set to 100 to provide stable anomaly score averaging across random trees without incurring excessive computational overhead. Increasing tree count beyond this range typically yields diminishing returns in anomaly detection stability. The maximum sample size was fixed at 512, balancing statistical representativeness and computational efficiency, consistent with standard Isolation Forest implementations. The contamination parameter was set to 0.05 to reflect an assumed upper-bound anomaly proportion within benchmark datasets while

preventing overfitting to rare outliers. This moderate contamination level enables sensitivity to anomalies without excessively inflating false positives.

The hybrid intrusion detection architecture combining Autoencoder feature extraction and Isolation Forest anomaly scoring is illustrated in Figure 2.

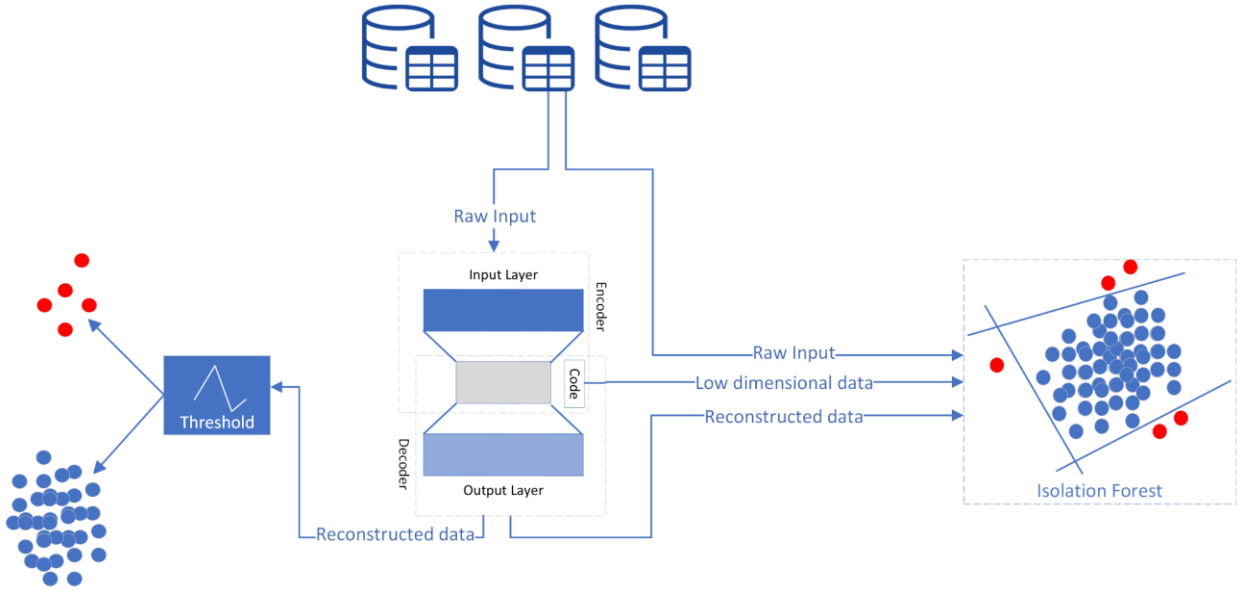

Figure 2: Hybrid Model Architecture (Almansoori, & Telek (2023))

Referring to Figure 2, the Autoencoder reduces the network data to a smaller feature space for reconstruction of normal traffic and submission to the Isolation Forest for anomaly scoring after transmitting the features and residuals to the Isolation Forest component. The Threshold component is used to filter residuals, while logical fusion or Hybrid OR/AND is applied to join AE and IF decisions in anomaly detection.

Hybrid Strategy and Evaluation Metrics

The hybrid model joined the AE and IF outputs by logical fusion rules to control precision and recall. There were two fusion methods:

1. **Hybrid OR:** An instance is classified as anomalous if flagged by either AE or IF. This configuration prioritises recall and detection sensitivity.
2. **Hybrid AND:** An instance is classified as anomalous only if both AE and IF agree. This configuration prioritises precision and reduced false alarm rates

The models were assessed using standard classification metrics based on the confusion matrix:

a. Accuracy: (TP+TN)/(TP+TN+FP+FN)

b. Precision: (TP/(TP+FP)) – the ratio of the detected anomalies that were actually anomalous.

c. Recall: (TP/(TP+FN)) – the ability to detect all attacks.

d. F1 Score: (2*Precision*Recall)/(Precision+Recall) – the harmonic mean of precision and recall.

Each of these values were calculated separately for the AE-only, IF-only, Hybrid OR, and Hybrid AND models for both datasets, which enabled the empirical assessment of generalization, trade-offs, and feasibility of real-time cybersecurity scenarios. All of the work was done in a Python programming environment with the assistance of the well-known tools TensorFlow, scikit-learn, and pandas.

4. RESULTS AND DISCUSSION

The proposed hybrid anomaly-based intrusion detection model was tested on the NSL-KDD and CICIDS2017 standard datasets to assess its performance, cross-generalisation, and influence of the logical fusion strategies. Four separate sturdy tests were conducted with four configurations: Autoencoder only (AE – only), Isolation Forest only (IF – only), Hybrid OR, and Hybrid AND. Performance metrics for the models included Accuracy, Precision, Recall, and F1-Score which allows comparing the models' respective sensitivities to detect true positives and yield false positives.

4.1 Experimental Results on NSL-KDD

Table 1 summarises the experiment results obtained on the NSL-KDD dataset

Table 1: Experimental Results on NSL-KDD

Model	Accuracy	Precision	Recall	F1-Score
AE-only	0.8300	0.9009	0.7880	0.8407

IF-only	0.6982	0.9183	0.5158	0.6605
Hybrid OR	0.8376	0.8975	0.8068	0.8498
Hybrid AND	0.6906	0.9247	0.4969	0.6465

The Autoencoder model demonstrated strong anomaly detection capability, achieving an F1-score of 0.8407 with recall of 78.8%. This means they can detect many attacks based on their performance in finding non-linear representations of normal network behaviour, as well as accurately detecting a large percentage of attack traffic, although they have a lower precision (compared to Isolation Forest (IF)). This indicates that the autoencoder is likely to produce false positives due to reconstruction errors being greater than the anomaly threshold.

Isolation Forest (IF) had the highest precision (0.9183), but had a low recall of 51.6%, meaning there were many undetected attacks. This is consistent with IF's statistical anomaly detection principle, because IF sometimes cannot identify an anomaly based on the underlying statistical relationship between featured components.

The best overall balance of the models was obtained through the Hybrid OR configuration, which achieved an F1-score of 0.8498 and a recall of 80.7%, suggesting that combining representation learning of an autoencoder and statistical scoring of isolation forests has improved overall detection coverage. Specifically, by using an OR fusion rule, combined with the differences in sensitivity of each model, IF will provide additional detection of anomalies over the use of an autoencoder alone, due to the limitations inherent in each model.

Conversely, the Hybrid AND configuration achieved the highest precision (0.9247) but at the cost of substantially reduced recall (49.7%). This behaviour reflects the restrictive nature of AND fusion, which requires both models to agree before classifying an instance as anomalous. While this reduces false alarms, it also increases the likelihood of missed detections.

4.2 Experimental Results on CICIDS2017

As indicated in Table 2, the same experimental configurations were applied to the CICIDS2017 dataset.

Table 2: Experimental Results on CICIDS2017

Model	Accuracy	Precision	Recall	F1-Score
AE-only	0.8342	0.8358	0.5794	0.6844
IF-only	0.7599	0.7520	0.3376	0.4660
Hybrid OR	0.8258	0.7665	0.6304	0.6918
Hybrid AND	0.7684	0.8963	0.2866	0.4343

The Autoencoder achieved moderate performance with an F1-score of 0.6844 and recall of 57.9%, suggesting that reconstruction-based detection remained effective but less stable under complex traffic distributions. Isolation Forest again exhibited conservative behaviour, achieving relatively low recall (33.8%), which indicates difficulty identifying anomalies within overlapping feature regions.

The Hybrid OR configuration again produced the strongest balanced performance, achieving an F1-score of 0.6918 with recall of 63.0%. This result indicates that decision-level fusion improves robustness under more complex datasets. By allowing anomaly detection from either model, the OR rule mitigates the limitations observed when using AE or IF independently.

In contrast, the Hybrid AND model produced the lowest recall (28.7%) despite achieving high precision. This confirms that strict agreement between models significantly reduces false alarms but also limits detection coverage in complex traffic environments.

4.3 Cross-Dataset Comparative Analysis

Table 3 – Cross-Dataset Comparative Analysis

Dataset	Model	Accuracy	Precision	Recall	F1-Score
NSL-KDD	Hybrid OR	0.8376	0.8975	0.8068	0.8498
CICIDS2017	Hybrid OR	0.8258	0.7665	0.6304	0.6918

As seen in Table 3 and Figure 3 below, the cross-validation results showed that, while the level of performance was higher in NSL-KDD, the generalization capability of the Hybrid OR model was good in both scenarios. The dominance of the F1 scores confirmed the position espoused by Mohammed and Telek (2023), stating that the best tradeoff between detection sensitivity and precision, which varies, can only be achieved in AE-IF hybrid models.

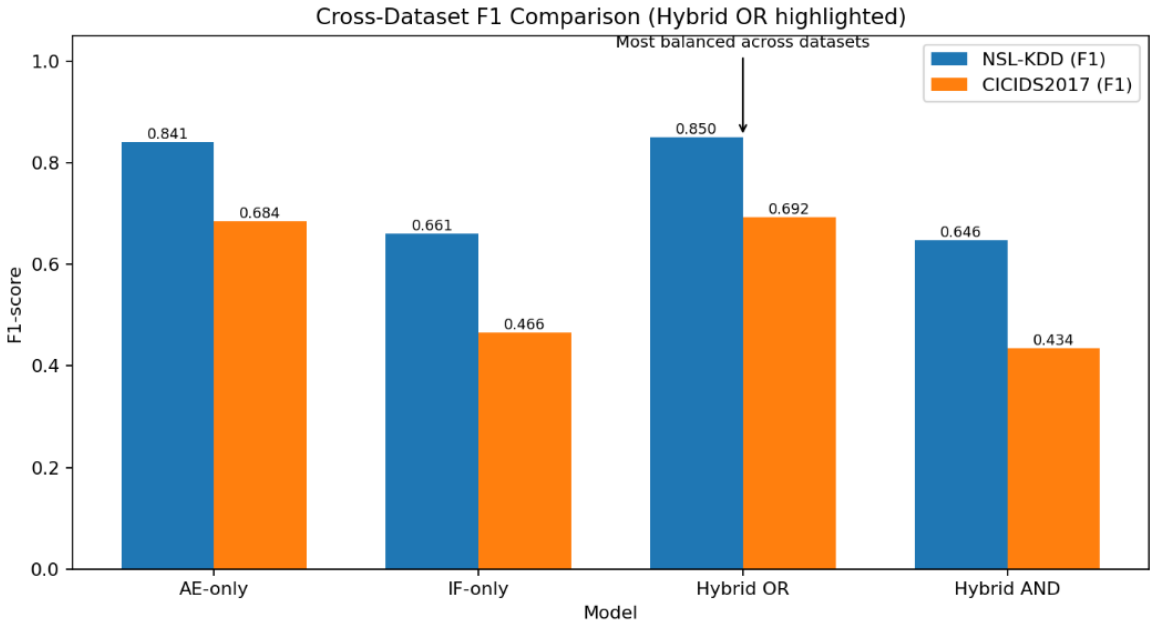

Figure 3: Cross-Dataset for F1 Comparison (Hybrid OR)

The Hybrid AND model's performance, in regards to winning minimal false alarms, did not align well with its performance in recall due to the restrictive threshold level used in detection. While this trend will be useful in situations when there is a relatively low tolerance of false alarms, it will likely not be appropriate in real-world network surveillance processes.

Discussion and Interpretation

The Hybrid OR configuration appeared to provide the best balance of detection performance among all models tested, as demonstrated by the complementary natures of both models. The Autoencoder identifies nonlinear relationships in network traffic through feature reconstruction, and the Isolation Forest identifies outliers in the observed data via statistical isolation. When these two models are combined using an OR fusion rule, anomalies detected by either model will remain, thus resulting in increased sensitivity when detecting anomalies compared to using either model independently.

The decreased performance when transitioning from the NSL-KDD dataset to the CICIDS2017 dataset reflects a greater level of complexity in today's datasets. Characteristics of CICIDS2017 include greater dimensionality and more heterogeneous traffic patterns. Thus, the challenge of differentiating between anomalous and normal traffic patterns is heightened. This is consistent with results from prior investigations conducted on this dataset (Engelen et al., 2021; Elsaid & Binbusayyis, 2024).

Although performance with the Hybrid OR configuration has been consistent between both datasets, both evaluations were conducted using benchmark datasets rather than real-time network traffic. Therefore, while there is an improved balance in terms of detection, there is still a need for subsequent testing under actual network conditions to verify whether this configuration can be utilized in operational environments.

Overall, the results suggest that combining representation learning and statistical isolation can improve the balance between precision and recall in anomaly detection tasks, particularly when compared with standalone Autoencoder or Isolation Forest models

5. CONCLUSION AND FUTURE WORK

This study proposed a hybrid anomaly-based intrusion detection system that integrates an Autoencoder (AE) and Isolation Forest (IF) within a sequential architecture. The approach combines deep representation learning with statistical isolation to enhance the detection of anomalous network behaviour. Experimental evaluation using the NSL-KDD and CICIDS2017

benchmark datasets demonstrated that the hybrid configuration improves the balance between detection sensitivity and false alarm control compared to standalone models. Among the evaluated configurations, the Hybrid OR decision rule produced the most balanced results, achieving F1-scores of 0.85 on NSL-KDD and 0.69 on CICIDS2017. These results indicate that decision-level fusion between reconstruction-based detection and statistical isolation can improve anomaly detection coverage while maintaining reasonable precision. The comparative analysis also showed that dataset characteristics significantly influence model behaviour, as performance declined on the more complex CICIDS2017 dataset due to higher dimensionality and traffic variability

Despite these encouraging findings, several limitations should be acknowledged. First, the experiments were conducted exclusively on benchmark datasets, which may not fully capture the complexity and variability of real-world network environments. Second, the evaluation was performed under offline experimental conditions rather than live network monitoring scenarios. Third, the hybrid model was tested using fixed hyperparameter settings and threshold configurations, and the influence of alternative parameter values on detection stability was not extensively explored.

Future research can extend this work in several directions. One important avenue involves evaluating the hybrid framework on live or streaming network traffic to assess operational feasibility in real-time monitoring environments. Additional studies may also investigate sensitivity analysis of hyperparameters, including reconstruction thresholds and contamination levels, to better understand their influence on anomaly detection stability. Expanding evaluation across multiple contemporary intrusion detection datasets would further strengthen cross-dataset generalisation analysis.

Further improvements may also be achieved by integrating explainable AI techniques to improve interpretability of anomaly detection decisions, particularly in operational cybersecurity environments. Investigating adversarial robustness and incorporating online learning mechanisms may also enhance adaptability to evolving attack patterns. Through these extensions, hybrid anomaly detection architectures may contribute to more robust and adaptive intrusion detection systems in future cybersecurity research.

6. REFERENCES

Abuabed, M., Alanazi, M., & Ameen, M. (2023). *STRIDE threat model-based framework for*

assessing the vulnerabilities of IoT systems. *Arabian Journal for Science and Engineering*, *48*(12), 16209–16224. https://doi.org/10.1007/s42452-024-06165-w

Adeniran, I., Adeniran, T., Familusi, O., & Adedayo, J. (2024). *Cybersecurity challenges in Nigeria's financial sector: Towards adaptive network defence models*. *Journal of Information Security and Digital Transformation, 6*(2), 88–103.

Ajagbe, S. A., Olawuyi, T. O., Odetunmibi, O. A., & Olanrewaju, A. M. (2024). *A comparison study of machine learning models using UNSW-NB15.* **SN Computer Science, 5**(3), [Article 03369]. https://doi.org/10.1007/s42979-024-03369-0

Ali, W. A., Manasa, K. N., Bendechache, M., Aljunaid, M. F., & Sandhya, P. (2023). *A review of current machine learning approaches for anomaly detection in network traffic*.

Almansoori, A., & Telek, M. (2023). *Intrusion detection based on Autoencoder and Isolation Forest in fog computing*. *International Journal of Communication Networks and Information Security, 15*(1), 12–20.

Alrayes, A., Alasmari, A., & Alshammari, M. (2024). *Intrusion detection in IoT systems using Denoising Autoencoder*. *IEEE Access, 12*, 54128–54139. https://doi.org/10.1109/ACCESS.2024.3382109

Chen, Y., Zhang, T., & Wang, Q. (2020). *A hybrid deep learning model for network intrusion detection system using Seq2Seq and ConvLSTM subnets*. *Procedia Computer Science, 177*, 150–158. https://doi.org/10.1016/j.procs.2020.10.023

Elsaid, M., & Binbusayyis, A. (2024). *A novel two-stage deep learning model for network intrusion detection (LSTM–AE)*. *Journal of Intelligent & Fuzzy Systems, 46*(3), 3779–3791. https://doi.org/10.3233/JIFS-240031

Engelen, J., & Papadopoulos, P. (2021). *Troubleshooting machine learning-based intrusion detection systems: Lessons from the CICIDS2017 dataset*. In *Proceedings of the Workshop on Traffic Measurement for Cybersecurity (WTMC 2021)*, 1–10.

Fuhnwi, D., Akbar, A., & Kim, D. (2023). *Hybrid machine learning approaches for anomaly-based intrusion detection in IoT environments*. *Sensors, 23*(15), 6593. https://doi.org/10.3390/s23156593

Hariharan, P., Anusha, S., & Srivastava, M. (2025). *A hybrid deep learning model for network intrusion detection system using Seq2Seq and ConvLSTM subnets*. *Multimedia Tools and Applications, 84*(7), 14725–14745.

Haque, M., & Soliman, M. (2025). *Transformer–Autoencoder with Isolation Forest and XGBoost*

for intrusion detection in wireless sensor networks. https://doi.org/10.20944/preprints202401.0185.v1

Hnamte, R., Singh, A., & Lalrinfela, C. (2023). *A novel two-stage deep learning model for*

network intrusion detection (LSTM-AE). Machine Learning and Security, 3(4), 205–214.

Kimanzi, G., Mutua, J., & Muriuki, S. (2024). *Digital finance, cybersecurity, and economic*

resilience: Emerging trends in Africa's fintech ecosystems. *African Journal of Cybersecurity, 5*(1), 32–47.

Kulkarni, S., Deshpande, P., & Gupta, R. (2022). *Hybrid Isolation Forest for anomaly detection*

in cybersecurity. *International Journal of Information Security Science, 11*(2), 55–66.

Labonne, A. (2020). *Anomaly-Based Network Intrusion Detection Using Machine Learning*

(PhD thesis).

Marteau, P., Ménard, T., & Auger, F. (2017). *Hybrid Isolation Forest: Enhancing anomaly*

detection sensitivity through distance-based scoring. *Pattern Recognition Letters, 105*, 16–23.

Moraboena, C., Singh, V., & Yadav, S. (2020). *A deep learning approach to network intrusion*

detection using Autoencoder. *International Journal of Recent Technology and Engineering*, *9*(6), 55–62.

Onyekpeze, C., Nnadi, E., & Alhassan, J. (2021). *Intelligent hybrid intrusion detection for IoT*

networks. *Nigerian Journal of Cybersecurity and Intelligence Systems, 3*(2), 70–83.

Prabu, K., Jeba, M., & Vinoth, R. (2023). *STRIDE-based cybersecurity threat modelling for*

IoT-enabled systems. *Materials Science and Engineering, 1299*(1), 012010. https://doi.org/10.1088/1757-899X/1299/1/012010

Sadaf, S., & Sultana, M. (2020). *Intrusion detection based on Autoencoder and Isolation Forest*

in fog computing. *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering, 8*(4), 1–7. https://doi.org/10.17148/IJIREEICE.2020.841

Talaei Khoei, A., & Kaabouch, N. (2023). *Anomaly detection in IoT networks using*

unsupervised machine learning algorithms. *Computers*, *12*(2), 34–47. https://doi.org/10.3390/computers12020034

Wang, V., Nnaji, H., & Jung, J. (2020). Internet Banking in Nigeria: Cyber Security Breaches, Practices and Capability.