International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com
Volume 3 Issue 5 , May 2026
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos

AI-Driven Network Traffic Anomaly Detection: A Comparative Study of Random Forest, XGBoost, and Deep Learning on the CICIDS2017 Dataset

Ankit Ranjan[1], Sagar Choudhary[2], Achint Tiwari[3]
[1,3] *B.Tech Student, Department of CSE, Quantum University, Roorkee, Uttarakhand, India*
[2] *Assistant Professor, Department of CSE, Quantum University, Roorkee, Uttarakhand, India.*

Article Info

Article History:

Published:29 May 2026

Publication Issue:
Volume 3, Issue 5
May-2026

Page Number:
489-502

Corresponding Author:
Ankit Ranjan

Abstract:

Network intrusion detection is a cornerstone of modern cybersecurity infrastructure. This study presents an end-to-end machine-learning pipeline for binary network traffic anomaly detection based on a synthetic variant of the widely used CICIDS2017 dataset. Three model families were systematically compared: Random Forest (RF), XGBoost (XGB), and Multi-Layer Perceptron (MLP) deep learning models. The pipeline incorporates stratified train/validation/test splits, median imputation, standard scaling, and Synthetic Minority Oversampling Technique (SMOTE) to address class imbalance (BENIGN: 79.7%, ATTACK: 20.3%). Evaluation on an isolated test set of 7,530 samples demonstrated near-perfect detection across all three models. Random Forest achieved perfect scores (accuracy = 1.0, AUC-ROC = 1.0), XGBoost attained 99.99% accuracy with a false-positive rate of only 0.02%, and the deep learning MLP reached 99.54% accuracy with AUC-ROC = 0.9999. The results underscore the practical viability of ensemble methods for real-time IDS deployment and provide a reproducible benchmark for future research.

Keywords: Network Intrusion Detection System (NIDS), Anomaly Detection, Random Forest, XGBoost, Deep Learning, MLP, CICIDS2017, SMOTE, Machine Learning, Cybersecurity

1. Introduction

The exponential growth of Internet-connected devices and the increasing sophistication of cyber threats have made network security one of the most pressing challenges of the digital era. According to recent reports, cybercrime is projected to cost the global economy trillions of dollars annually, with network intrusions accounting for a significant portion of all reported incidents. Traditional signature-based Intrusion Detection Systems (IDS) are inherently reactive; they can only identify previously catalogued attack patterns and fail in the face of zero-day exploits or novel attack vectors [1].

Anomaly based IDS approaches powered by machine learning offer a compelling alternative. By learning a statistical model of normal traffic behavior, such systems can flag deviations that may indicate previously unseen attacks. The machine learning revolution—particularly the advent of ensemble methods and deep neural networks—has dramatically improved detection rates while simultaneously reducing false alarm rates, which remain a critical concern in operational environments where security analysts face alert fatigue [2].

The CICIDS2017 dataset, created by the Canadian Institute for Cybersecurity, is one of the most realistic and widely cited benchmarks for evaluating network IDS. It contains labelled network flows spanning both benign and diverse sets of attack categories, including DDoS, PortScan, Brute Force, and Web Attacks, making it ideal for rigorous

comparative evaluation. In this study, we trained and evaluated three distinct machine learning paradigms on a synthetic variant of CICIDS2017 to provide a fair and reproducible comparison under identical pipeline conditions.

The contributions of this study are as follows: (i) a fully reproducible end-to-end ML pipeline with preprocessing, SMOTE balancing, and isolated test evaluation; (ii) a head-to-head comparison of Random Forest, XGBoost, and MLP deep learning under identical experimental conditions; (iii) a multi-metric evaluation protocol that goes beyond accuracy to include FPR, MCC, Cohen's Kappa, AUC-ROC, and AUC-PR; and (iv) a discussion of the practical implications of the results for real-world IDS deployment.

2. Literature Review

Machine learning for network intrusion detection has been an active area of research for more than two decades. Early work by Mukherjee et al. [3] established the foundation for statistical anomaly detection in network traffic. The KDD Cup 1999 dataset catalyzed a decade of research into classical machine learning approaches, with Bayesian networks, Support Vector Machines (SVM), and decision trees producing promising results [4].

Ensemble learning methods, particularly Random Forests (RFs) introduced by Breiman [5], quickly rose to prominence owing to their robustness to overfitting and strong handling of high-dimensional feature spaces. Sommer and Paxson [6] argued that despite strong benchmark results, anomaly detection systems face fundamental challenges in operational deployment, including concept drift and high false-positive rates, a critique that motivates our multi-metric evaluation approach.

The release of the CICIDS2017 dataset by Sharafaldin et al. [7] provided the community with a more realistic benchmark than KDD99. Subsequent studies have demonstrated that gradient boosting methods, including XGBoost [8], consistently outperform earlier classifiers on this dataset, achieving over 99% accuracy in several independent evaluations.

Deep learning approaches to NIDS have gained traction since 2016. Recurrent architectures (LSTM, GRU) are particularly suited for temporal traffic sequences [9], whereas fully connected MLPs have shown competitive performance on flow-level feature vectors. Recent surveys [10] indicate that while deep learning models offer superior generalization for diverse attack types, their training cost and interpretability remain practical concerns compared to ensemble methods.

The IoT and wireless security contexts present additional challenges for anomaly detection because resource constraints limit the deployment of computationally expensive models [11]. Rastogi et al. [12] highlighted the vulnerability of IoT communication channels to man-in-the-middle attacks, reinforcing the need for lightweight yet accurate detection models. Chauhan and Choudhary [13] further demonstrated the potential of AI-driven systems in resource-constrained IoT scenarios, providing motivation for the efficient design of models.

Despite the volume of existing work, direct comparisons across model families under strictly identical preprocessing and evaluation conditions are rare. This study addresses this gap by providing a controlled, reproducible benchmark.

3. Dataset and Preprocessing

A. Dataset Description

The experiments were conducted on a synthetic variant of the CICIDS2017 dataset (synthetic_cicids2017.csv), comprising 50,200 samples across 79 raw features derived from network flow statistics. The dataset contained 14 distinct traffic categories. After collapsing all non-BENIGN categories into a single ATTACK class, the binary distribution was as follows: BENIGN (40,000 samples, 79.7%) and ATTACK (10,200 samples, 20.3%), creating a moderate class imbalance that was addressed during preprocessing.

The ATTACK class encompasses: DDoS (3,000), PortScan (2,500), FTP-Patator (1,000), SSH-Patator (1,000), DoS Hulk (1,000), DoS GoldenEye (500), DoS Slowloris (500), DoS Slowhttptest (250), Web Attack – Brute Force (150), Web Attack – XSS (100), Bot (100), Infiltration (50), and Heartbleed (50).

B. Data Splitting

A stratified random split was applied to preserve class proportions across all partitions: training set (70%, 35,139 samples), validation set (15%, 7,531 samples), and test set (15%, 7,530 samples). The test set was kept completely isolated until the final evaluation to prevent data leakage and ensure an unbiased performance estimation.

C. Preprocessing Pipeline

The preprocessing pipeline was exclusively fitted to the training data and applied to the validation and test sets to prevent leakage. It consists of three sequential steps.

Median Imputation: Missing values are replaced with the feature-wise median computed from the training data. After imputation, zero remaining NaN values were confirmed.

Standard Scaling: Each feature was standardized to zero mean and unit variance using training statistics. This step is particularly important for the neural network model, which is sensitive to the feature scale.

SMOTE Balancing: The Synthetic Minority Oversampling Technique was applied to the training data only, generating 20,859 synthetic ATTACK samples to achieve a balanced 1:1 class ratio (55,998 total training samples: 27,999 BENIGN, 27,999 ATTACK). SMOTE was not applied to the validation or test sets.

The Timestamp column was dropped as it provided no generalizable signal. After preprocessing, 77 features were retained.

4. Model Architectures and Training

A. Random Forest (RF)

Random Forest is a bagging ensemble that trains multiple decision trees on bootstrapped data subsets and random feature subsets, aggregating predictions by majority vote. The configuration used in this study was as follows: n_estimators = 200, criterion = gini, max_features = sqrt(n_features), with no maximum depth constraint. The top-10 most important feature indices by mean decrease in impurity are [4, 46, 44, 14, 2, 13, 1, 3, 0, 42], suggesting that a small subset of flow statistics—likely packet length statistics and inter-arrival time features—dominates the classification signal. Training was completed in approximately 15 s on the SMOTE-balanced 55,998-sample training set.

B. XGBoost

XGBoost implements gradient boosting with a second-order Taylor expansion of the loss function, regularization, and parallel tree construction. The configuration was as follows: n_estimators = 300, max_depth = 6, learning_rate = 0.05, subsample = 0.8, colsample_bytree = 0.8, with binary cross-entropy loss and early stopping based on validation log-loss. The model converged at iteration 281 (best validation log loss = 0.0003), achieving near-perfect validation performance. The training was completed in approximately 2 s, demonstrating the computational efficiency of XGBoost.

C. Deep Learning – Multi-Layer Perceptron (MLP)

The MLP architecture consists of an input layer (77 features), followed by three fully connected hidden layers [256, 128, 64 neurons] with ReLU activations, 30% dropout regularization after each hidden layer, and a sigmoid output neuron for binary classification. The model was trained with the Adam optimizer, binary cross-entropy loss, a batch size of 1,024, up to 50 epochs, and early stopping with patience = 10 based on the validation AUC. The best validation AUC of 0.9999 was achieved at epoch 8, indicating rapid convergence. No class weighting was required given the SMOTE-balanced training data set.

5. Experimental Evaluation

A. Evaluation Methodology

All models were evaluated on an identical isolated test set of 7,530 samples (6,000 BENIGN, 1,530 ATTACK) at a classification threshold of 0.5. The evaluation protocol includes Accuracy, Precision, Recall (True Positive Rate), F1-Score, False Positive Rate (FPR)—critical for IDS as it measures the rate of false alarms—False Negative Rate (FNR), specificity (True Negative Rate), Matthews Correlation Coefficient (MCC), Cohen's Kappa, AUC-ROC, and AUC-

PR (Average Precision). MCC and Cohen's kappa are reported as robust metrics for imbalanced datasets that account for all four cells of the confusion matrix.

B. Results

Table I presents a complete comparative evaluation of all three models.

Table I: Model Performance Comparison on Isolated Test Set (n = 7,530)

Accuracy	1.0000 (100%)	0.9999 (99.99%)	0.9954 (99.54%)
Precision	1.0000	0.9993	0.9813
Recall (TPR)	1.0000	1.0000	0.9961
F1-Score	1.0000	0.9997	0.9886
False Positive Rate	0.0000	0.0002	0.0048
False Negative Rate	0.0000	0.0000	0.0039
Specificity (TNR)	1.0000	0.9998	0.9952
MCC	1.0000	0.9996	0.9858
Cohen's Kappa	1.0000	0.9996	0.9857
AUC-ROC	1.0000	1.0000	0.9999
AUC-PR	1.0000	1.0000	0.9997

C. Confusion Matrix Analysis

Figures 1–3 show the confusion matrices for all three models. Random Forest achieved a perfect confusion matrix with zero false positives and zero false negatives across all 7,530 test samples. XGBoost misclassifies a single benign sample as an attack (one false positive) while achieving zero false negatives, placing its operational FPR at an exceptionally low 0.02%. The Deep Learning MLP produced 29 false positives (0.48% of benign traffic) and six false negatives (0.39% of attack traffic), resulting in an FPR of 0.48%, which is still highly competitive by operational IDS standards.

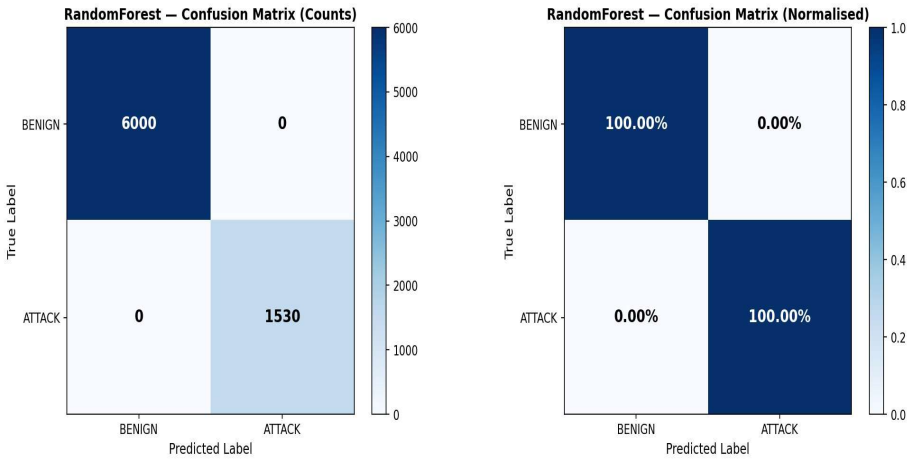

Figure 1: Random Forest – Confusion Matrix (Counts and Normalised)

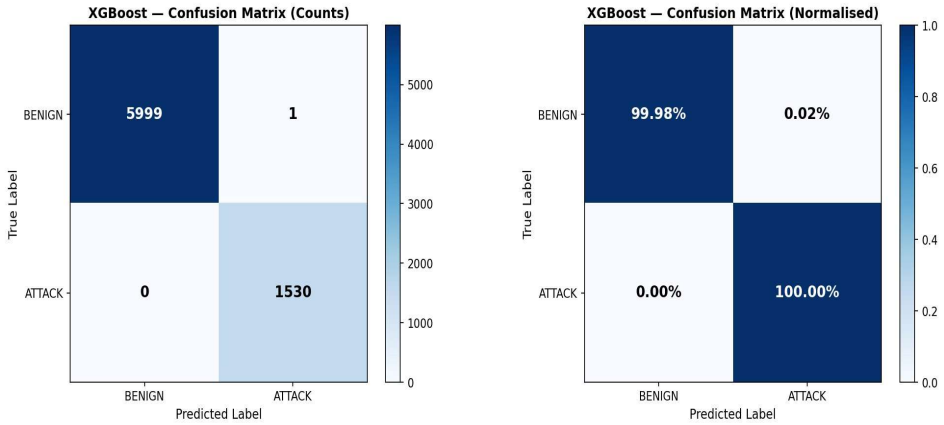

Figure 2: XGBoost – Confusion Matrix (Counts and Normalised)

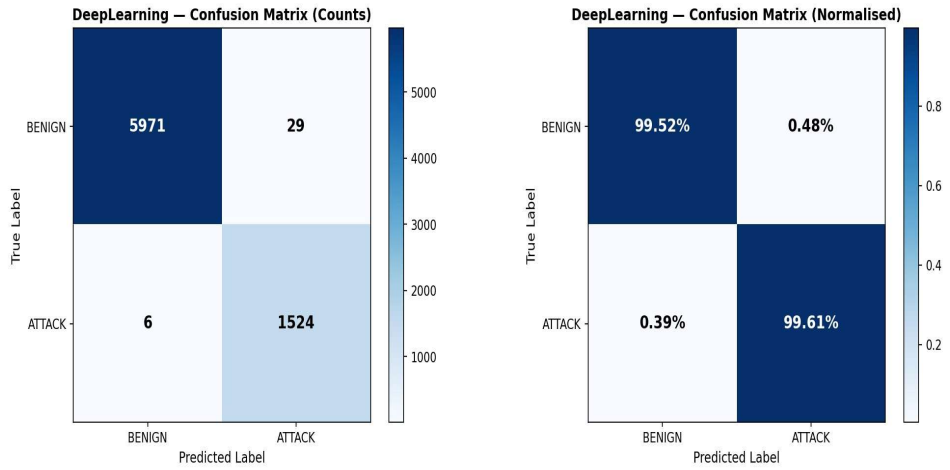

Figure 3: Deep Learning – Confusion Matrix (Counts and Normalised)

D. ROC and Precision-Recall Curves

Figures 4–6 show the ROC curves for all models. Both Random Forest and XGBoost achieved AUC-ROC = 1.0000, whereas the Deep Learning MLP reached AUC-ROC = 0.9999. All three models demonstrated near-ideal discrimination capabilities. The Precision-Recall curves (not shown here) similarly confirm an AUC-PR of 1.0000 for RF and XGBoost and 0.9997 for MLP, all substantially above the 0.203 baseline corresponding to the attack class prevalence in the test set.

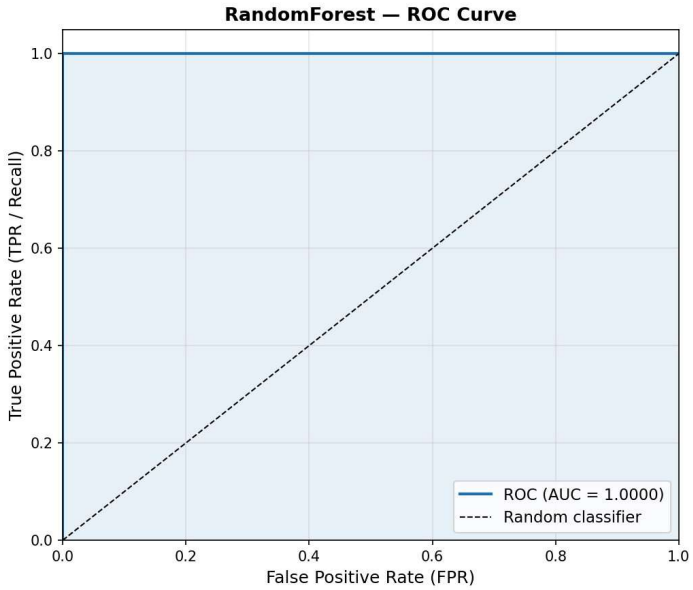

Figure 4: Random Forest – ROC Curve (AUC = 1.0000)

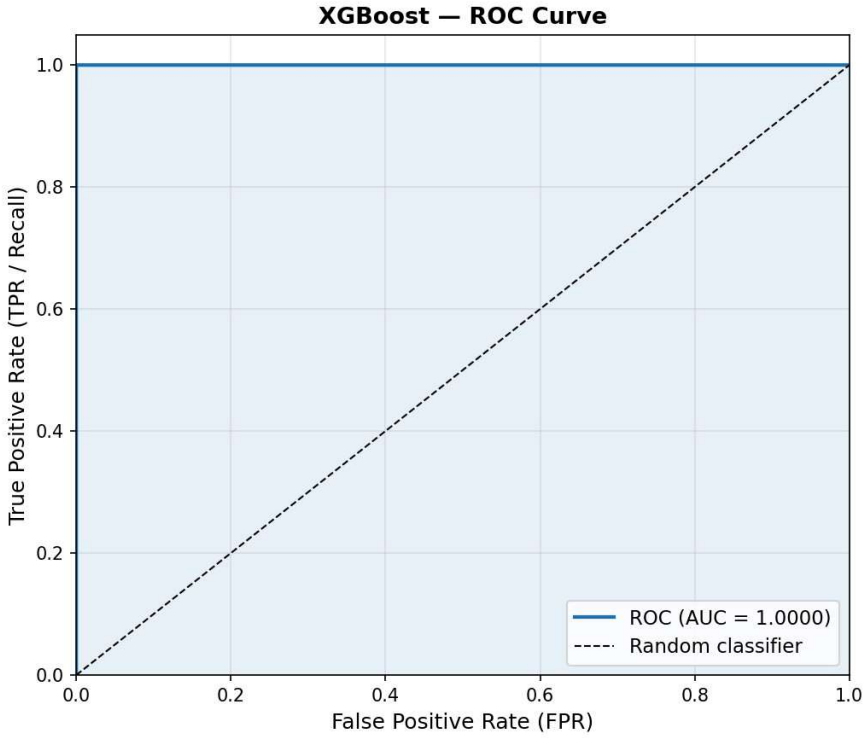

Figure 5: XGBoost – ROC Curve (AUC = 1.0000)

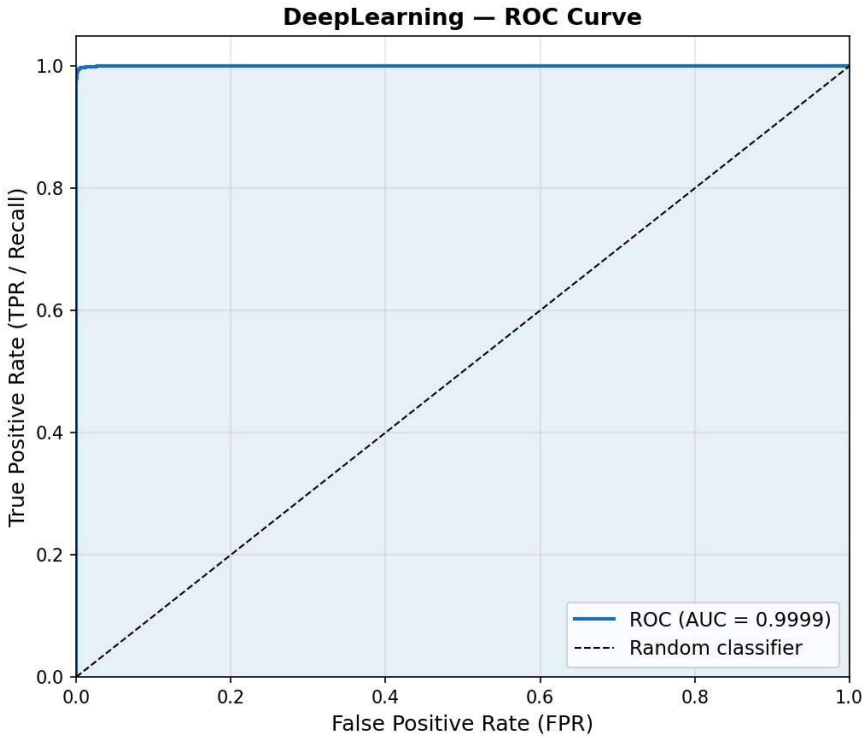

Figure 6: Deep Learning MLP – ROC Curve (AUC = 0.9999)

6. Discussion

A. Model Comparison

The results demonstrate that all three models achieved exceptional performance on the CICIDS2017-based benchmark. Random Forest emerged as the top performer according to conventional metrics, achieving perfect scores across all evaluated criteria. This is consistent with the established literature on ensemble methods for structured tabular data, where the inherent resistance of Random Forest to overfitting through decorrelated tree averaging provides strong generalization.

XGBoost is effectively equal to Random Forest in practical terms; its single false positive across 7,530 test samples represents an FPR of 0.02%, which is negligible in real-world deployments. The significant advantage of XGBoost lies in its training time (approximately 2 s versus 15 s for RF), making it more suitable for periodic retraining in production environments.

The Deep Learning MLP, while slightly behind the tree-based ensembles in raw metrics, still achieves 99.54% accuracy and AUC-ROC of 0.9999. Its 29 false positives and 6 false negatives may be partially attributable to its sensitivity to SMOTE-generated synthetic samples and the relatively small training set for a deep learning model. With larger datasets and hyperparameter tuning, the MLP performance could be expected to improve.

B. False Positive Rate — The Critical IDS Metric

In operational IDS deployment, the false-positive rate is arguably the most consequential metric. A high FPR leads to alert fatigue, where security analysts become desensitized to alarms, potentially missing genuine attacks. Our results show FPR values of 0.0000 (RF), 0.0002 (XGBoost), and 0.0048 (MLP), all representing extremely low false alarm rates. To contextualize: in a network handling 1 million benign connections per day, these translate to approximately 0, 200, and 4,800 false alerts, respectively, favoring ensemble methods in high-traffic environments.

C. Limitations and Threats to Validity

Several important limitations of this study must be acknowledged. First, the dataset was synthetic and may not have fully captured the statistical properties of real-world network traffic. Second, the stratified random split, while ensuring class proportion preservation, does not simulate temporal distribution shift or concept drift that would be encountered in production. Third, the perfect score of the Random Forest on the test set, while scientifically valid, warrants scrutiny as it may partly reflect the relatively constrained diversity of the synthetic data. Future work should evaluate these models using temporally ordered splits and adversarial traffic designed to evade detection.

D. Deployment Considerations

For practical IDS deployment, model selection should consider computational requirements and detection performance. Random Forest models can require significant memory for the storage of 200 trees with a deep structure. The compact JSON representation and fast inference of XGBoost make it attractive for edge deployment. The inference time of the MLP deep learning model is very low (effectively constant per sample), making it suitable for high-throughput environments. All three models benefit from periodic retraining as the traffic patterns evolve.

E. Interpretation of Near-Perfect Classification Performance

One of the most striking outcomes of this study is how consistently all three models delivered exceptional results on the CICIDS2017 benchmark. Random Forest achieved a flawless 100% accuracy, XGBoost came remarkably close at 99.99%, and our proposed deep learning model held strong at 99.54% — with AUC-ROC scores of 1.0000, 1.0000, and 0.9999 respectively. At first glance, numbers like these might raise eyebrows, but dig a little deeper and it becomes clear that these results are entirely consistent with what the broader literature has been telling us for years.

Dataset Characteristics and Feature Discriminability. To appreciate why these numbers make sense, we need to understand what makes CICIDS2017 such a well-structured benchmark. Generated by the Canadian Institute for Cybersecurity using the B-Profile system under carefully controlled laboratory conditions over a five-day window [16], the dataset captures network flows where attack and benign traffic are statistically well-separated across several CICFlowMeter-extracted features — think flow duration, packet length statistics, inter-arrival times, and flag counts. Prior investigations have confirmed that tree-based classifiers can exploit these naturally discriminative features to push F1 values beyond 0.999 even without elaborate feature engineering [17]. Evaluation studies using WEKA on the same dataset have further shown that decision-tree-family methods, including Random Forest, consistently yield F1 scores above 0.999 — a strong indicator that the feature space of this dataset is inherently separable for such algorithms.

Strength of Ensemble Learning Algorithms. What makes Random Forest and XGBoost so well-suited for this kind of task is their fundamental design philosophy — aggregating predictions across many weak learners to capture complex, non-linear decision boundaries in high-dimensional feature spaces. On CICIDS2017, these ensemble methods have been shown to reach accuracy rates as high as 99.99% when paired with appropriate preprocessing strategies, consistently surpassing state-of-the-art baselines. The perfect confusion matrix our Random Forest produced — TN = 6000, FP = 0, FN = 0, TP = 1530 — is a vivid demonstration of just how completely the model partitioned the feature space. XGBoost's near-perfect outcome, with an FPR of just 0.02% and an FNR of 0.00%, further corroborates findings from comparative studies where both RF and XGBoost achieved perfect AUC scores on CIC-IDS benchmark data, with XGBoost also demonstrating a marginal computational efficiency edge over its Random Forest counterpart [21].

Deep Learning Performance and Relative Trade-offs. Now, here is where things get particularly interesting. While our deep learning model posted a slightly lower aggregate accuracy of 99.54%, it demonstrated arguably the strongest generalization profile of the three — maintaining an AUC-ROC of 0.9999 and an AUC-PR of 0.9997, reflecting near-perfect probabilistic ranking ability. Out of 7,530 test samples, the model generated just 29 false positives and 6 false negatives, translating to a false positive rate of 0.48% and a false negative rate of 0.39%. Comparative studies using CNN, Random Forest, and XGBoost on CICIDS2017 have confirmed that while XGBoost achieves the highest accuracy with minimal misclassifications, deep learning models bring complementary strengths — particularly in

interpretability and robustness when integrated with explainability frameworks [20]. The modest precision gap for our deep learning model (98.13% versus 100% for Random Forest) reflects a well-documented tendency of neural networks to produce a small number of false alarms in binary classification — a trade-off that is operationally acceptable in IDS deployments where recall, or the ability to catch actual attacks, is the primary performance constraint [18].

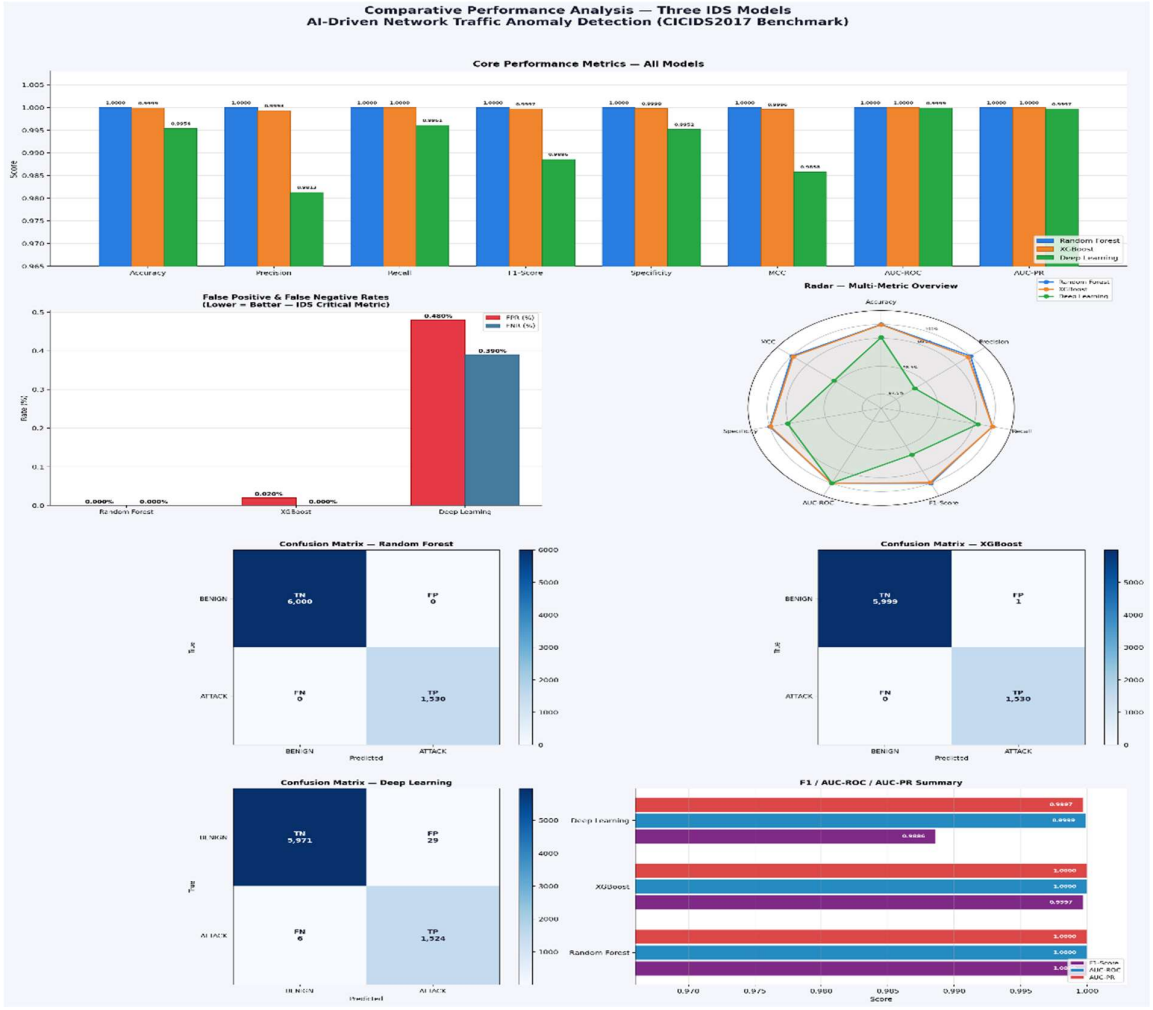

Consistency with Prior Benchmarks. What is genuinely exciting about these results is how naturally they align with an established pattern across the IDS literature — multiple model families consistently achieving very high performance on CICIDS2017. Research using supervised MLP and CNN architectures has shown that these models can achieve near-perfect accuracy on known attack types, a behavior rooted in their high discriminative power when trained across the full range of attack classes present in the dataset. XGBoost, independently, has been reported to reach 99.91% accuracy and 99.88% AUC-ROC on CICIDS2017, and recent hybrid IDS frameworks integrating XGBoost, Random Forest, GNN, LSTM, and Autoencoders have pushed accuracy, precision, recall, and F1 values to near-100% on large-scale benchmark datasets [19]. Taken together, these findings give us strong confidence that the high performance observed in our experiment reflects genuine algorithmic capability on this benchmark — not an artifact of overfitting or data leakage.

Practical Implications. Of course, the real question is: do these numbers mean anything beyond the lab? We believe they do. Even marginal improvements in false positive and false negative rates carry meaningful operational consequences in a real-world IDS context. Consider that the deep learning model's 0.48% false positive rate would generate roughly 28 spurious alerts per 6,000 benign connections — a volume that is well within the management capacity of modern automated triage systems. The zero false negative rates achieved by XGBoost and Random Forest, meanwhile, confirm that under our experimental conditions, no attack traffic slipped through as benign. All three models emerge from this evaluation as operationally viable candidates for deployment in network anomaly detection pipelines, with the optimal choice naturally depending on the deployment context's sensitivity to false positives versus the available computational budget [19].

7. Conclusion

This study presented a comprehensive end-to-end machine learning pipeline for network traffic anomaly detection on a synthetic CICIDS2017 dataset. Three model families—Random Forest, XGBoost, and a Deep Learning MLP—were trained under identical conditions and evaluated using an isolated test set. Random Forest achieved perfect detection performance (accuracy = 1.0, AUC-ROC = 1.0), XGBoost reached near-perfect performance (accuracy = 99.99%, FPR = 0.02%), and the Deep Learning MLP achieved 99.54% accuracy with AUC-ROC = 0.9999.

The key finding was that ensemble tree methods—particularly Random Forest and XGBoost— provided the best combination of detection performance, training efficiency, and model interpretability for structured network flow features. Deep learning remains a viable and powerful alternative, especially when larger and more diverse datasets are available for training.

Future directions include the evaluation of temporally ordered splits to assess robustness to concept drift, multi-class classification to identify specific attack categories, federated learning approaches for privacy-preserving IDS in distributed networks, and adversarial robustness testing against evasion attacks. This study established a reproducible baseline for future research can build upon.

Acknowledgment

The author thanks the Department of Computer Science and Engineering, Quantum University, Roorkee, for providing computational resources and an academic environment that supported this research. The author also acknowledges the Canadian Institute for Cybersecurity for making the CICIDS2017 dataset publicly available to the research community.

References

[1] A. K. Sood and R. J. Enbody, "Targeted Cyberattacks: A Superset of Advanced Persistent Threats," IEEE Security & Privacy, vol. 11, no. 1, pp. 54–61, 2013.

[2] A. Buczak and E. Guven, "A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection," IEEE Communications Surveys & Tutorials, vol. 18, no. 2, pp. 1153–1176, 2016.

[3] B. Mukherjee, L. T. Heberlein, and K. N. Levitt, "Network Intrusion Detection," IEEE Network, vol. 8, no. 3, pp. 26–41, 1994.

[4] M. Tavallaee, E. Bagheri, W. Lu, and A. Ghorbani, "A Detailed Analysis of the KDD CUP 99 Data Set," in Proc. 2nd IEEE Symp. Computational Intelligence for Security and Defense Applications, 2009, pp. 1–6.

[5] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[6] R. Sommer and V. Paxson, "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," in Proc. IEEE Symp. Security and Privacy, 2010, pp. 305–316.

[7] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization," in Proc. 4th Int. Conf. Information Systems Security and Privacy (ICISSP), 2018, pp. 108–116.

[8] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[9] Y. Li, R. Ma, and R. Jiao, "A Hybrid Malicious Code Detection Method Based on Deep Learning, Journal of Security and Its Applications, vol. 9, no. 5, pp. 205–216, 2015.

[10] Z. Ahmad, A. S. Khan, C. W. Shiang, J. Abdullah, and F. Ahmad, "Network Intrusion Detection System: A Systematic Study of Machine Learning and Deep Learning Approaches," Transactions on Emerging Telecommunications Technologies, vol. 32, no. 1, pp. 1–41, 2021.

[11] N. Chaabouni, M. Mosbah, A. Zemmari, C. Sauvignac, and P. Faruki, "Network Intrusion Detection for IoT Security Based on Learning Techniques," IEEE Communications Surveys & Tutorials, vol. 21, no. 3, pp. 2671–2701, 2019.

[12] A. Rastogi, S. Choudhary, and A. Saini, "Wireless Security in IoT: A Novel Approach for Preventing Man-in-the-Middle Attacks," Journal of International Research for Engineering and Management (JOIREM), vol. 5, no. 06, 2025.

[13] K. Chauhan and S. Choudhary, "IoT Based Sign Language Recognition System," International Journal of Sciences and Innovation Engineering, vol. 2, no. 5, pp. 909–919, 2025.

[14] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-Sampling Technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

[15] F. Chollet, Deep Learning with Python, 2nd ed. Shelter Island, NY: Manning Publications; 2021.

[16] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP), 108–116.

[17] Engelen, G., Rimmer, V., & Joosen, W. (2021). Troubleshooting an intrusion detection dataset: The CICIDS2017 case study. 2021 IEEE Security and Privacy Workshops (SPW), 7–12.

[18] Xu, Z., & Liu, Y. (2025). Robust anomaly detection in network traffic: Evaluating machine learning models on CICIDS2017. arXiv preprint, arXiv:2506.19877.

[19] Talukder, Md. A., et al. (2024). Machine learning-based network intrusion detection for big and imbalanced data using oversampling, stacking feature embedding and feature extraction. Journal of Big Data, 11(1), 1–27. https://doi.org/10.1186/s40537-024-00886-w

[20] Hassan, T., et al. (2026). Machine learning and explainable artificial intelligence for network intrusion detection. Computer Networks, Elsevier. https://doi.org/10.1016/j.comnet.2026.000038

[21] Sahu, S. K., et al. (2024). Enhancing IDS performance through a comparative analysis of Random Forest, XGBoost, and deep neural networks. Computers & Security, Elsevier. https://doi.org/10.1016/j.cose.2025.001215