Hybrid ANN–XGBoost Framework for Intelligent Intrusion Detection in Internet of Medical Things Networks

Okeke, O. C.[1], Mordi, S. I.[2], Mgbeafulike, I. J.[3], Nwaoha, S. O.[4]

[1,2,3]Department of Computer Science, Chukwuemeka Odumegwu Ojukwu University, P.M.B. 02, Uli, Anambra State, Nigeria.
[4]Department of Software Engineering, Nnamdi Azikiwe University, Awka, Anambra State,
*Corresponding Author: drnwaohastephen@gmail.com

ABSTRACT

The Internet of Medical Things (IoMT) connects medical devices, sensors, monitoring platforms, and healthcare information systems, enabling continuous data exchange and remote healthcare services. However, increased connectivity expands the cybersecurity attack surface of healthcare infrastructures and exposes sensitive information and services to potential threats. This study presents a hybrid Artificial Neural Network–Extreme Gradient Boosting (ANN–XGBoost) framework for intelligent intrusion detection in IoMT environments. The proposed approach combines the nonlinear learning capability of Artificial Neural Network (ANN) with the ensemble learning capability of XGBoost through a weighted probability averaging mechanism. An experimental machine-learning workflow involving data acquisition, preprocessing, feature engineering, model development, hybrid integration, and comparative evaluation was adopted, together with an Agile-oriented software-development process. Logistic Regression, Random Forest, ANN, and standalone XGBoost were evaluated alongside the proposed hybrid model. Based on the experimental results, Logistic Regression, Random Forest, ANN, and XGBoost achieved accuracies of 95.4%, 97.6%, 97.1%, and 97.8%, respectively, while the hybrid ANN–XGBoost model achieved an accuracy of 98.6%, precision of 0.983, recall of 0.989, and F1-score of 0.986. The hybrid model also achieved a ROC-AUC of 0.99 and an average prediction response time of 1.18 seconds. A web-based implementation using Python, Flask, machine-learning libraries, and SQLite3 was developed to support authentication, data upload, preprocessing, prediction, and result logging. The study demonstrates a practical framework for integrating neural and gradient-boosting models for intelligent intrusion detection in IoMT environments.

Keywords: Internet of Medical Things; Intrusion Detection; Artificial Neural Network; XGBoost; Machine Learning; Healthcare Cybersecurity; Hybrid Learning; ROC-AUC.

I. INTRODUCTION

The Internet of Medical Things (IoMT) represents the integration of Internet-connected technologies with medical devices, sensors, healthcare information systems, and communication infrastructures. IoMT technologies facilitate continuous patient monitoring, remote healthcare services, automated data collection, and communication between patients, healthcare professionals, and connected medical equipment. These capabilities have contributed to the transformation of healthcare delivery by enabling medical information to be generated, processed, and exchanged across distributed environments (Naghib et al., 2025; Alalhareth & Hong, 2023). Despite these benefits, the increasing interconnection of medical devices introduces substantial cybersecurity concerns. Healthcare environments contain highly sensitive information and increasingly depend on network-connected devices for monitoring and service

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

delivery. Consequently, a successful cyberattack may affect both information assets and the availability, integrity, or reliability of healthcare services. Recent IoMT security research emphasizes that device heterogeneity, limited computational resources, privacy requirements, interoperability challenges, and evolving cyber threats create significant difficulties for securing connected healthcare environments (Naghib et al., 2025; Alalhareth & Hong, 2023).

Intrusion Detection Systems (IDSs) constitute an important layer of cybersecurity for identifying malicious activities within networked environments (Nwaoha & Okeke, 2025). An IDS continuously monitors network activities and analyzes observed traffic or system behaviour to identify suspicious or malicious events and generate appropriate alerts (Nwaoha & Mmaduekwe, 2026; Naghib et al., 2025). Conventional signature-based IDSs identify threats by matching observed behaviour against previously known attack signatures. Although this approach can be effective for known attacks, its dependence on predefined signatures can limit its ability to identify previously unseen or modified threats. Anomaly-based approaches, in contrast, attempt to identify deviations from established patterns of normal behaviour; however, changes in network conditions and legitimate traffic behaviour may contribute to false-positive classifications (Naghib et al., 2025).

These limitations have encouraged the application of machine learning and deep learning techniques to intrusion detection in healthcare and IoMT environments. Machine-learning models can learn relationships between network characteristics and attack classes, thereby reducing dependence on manually defined detection rules. Deep learning models can further learn complex nonlinear representations from high-dimensional data and have therefore been investigated for detecting sophisticated cyber threats. Recent studies have examined Artificial Neural Networks (ANN), Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), Random Forest, gradient boosting, and other intelligent approaches for IoMT and healthcare intrusion detection (Kumaar et al., 2022; Alalhareth & Hong, 2023; Naghib et al., 2025). The objectives are to:

(a). design a machine-learning framework for identifying legitimate and malicious activities in IoMT network traffic;
(b). implement Artificial Neural Network (ANN) and Extreme Gradient Boosting (XGBoost) models for network intrusion classification; and (c) integrate ANN and XGBoost using a weighted probability-averaging mechanism to form the proposed hybrid intrusion-detection framework;
(c). compare the performance of the ANN, XGBoost, and hybrid ANN–XGBoost models with selected baseline classifiers;
(d). evaluate the model's using accuracy, precision, recall, F1-score, ROC-AUC, confusion matrix analysis, and prediction response time;
(e). deploy the proposed framework in a web-based environment supporting authentication, dataset upload, preprocessing, intrusion prediction, and prediction-result logging.

II. CONCEPTUAL REVIEW

A. Cybersecurity in Internet of Medical Things

IoMT environments consist of interconnected medical devices, sensors, gateways, healthcare information systems and communication infrastructures. Examples include wearable monitoring devices, medical sensors, remote patient-monitoring systems and connected medical equipment. The interconnected structure of IoMT introduces multiple potential attack surfaces. Security threats may involve unauthorized access, denial-of-service attacks, malware, spoofing, replay attacks, man-in-the-middle attacks and other forms of malicious network activity. The security requirements of IoMT extend beyond conventional information security because cyber incidents can potentially affect healthcare operations. Confidentiality is required to protect patient

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

information, integrity is necessary to prevent unauthorized modification of medical data, and availability is necessary to ensure that connected healthcare services remain operational. Recent surveys emphasize that AI-based intrusion detection has become an important research direction for addressing these challenges. Naghib et al. (2025), in a systematic review of IoMT intrusion detection, categorized existing studies according to AI techniques, datasets, security requirements, intrusion-detection processes and evaluation metrics. The authors also identified dataset and deployment challenges that remain open research issues. A recent IEEE Access survey similarly emphasizes security and privacy challenges across IoMT architecture layers and discusses the application of machine learning, deep learning and blockchain for threat detection and secure data transmission.

B. Machine Learning-Based Intrusion Detection

Machine learning transforms intrusion detection into a data-driven classification or anomaly detection problem. Given labelled network traffic, supervised algorithms learn relationships between traffic features and attack classes. Logistic Regression provides a relatively simple baseline and is useful where relationships between predictors and the classification outcome can be adequately represented by a linear decision boundary. Random Forest combines multiple decision trees and is capable of modelling nonlinear relationships. Its ensemble structure can improve robustness compared with individual decision trees. XGBoost uses gradient boosting to construct an ensemble of decision trees sequentially. Each successive tree contributes to correcting errors from previous trees. This makes XGBoost particularly useful for structured tabular data. Kulshrestha and Vijay, (2024) evaluated several machine-learning methods for IoMT intrusion detection and reported strong performance from ensemble methods. Their study included Logistic Regression, Random Forest, Adaptive Boosting, Gradient Boosting and XGBoost. (IDEAS/RePEc). However, machine-learning performance depends strongly on feature representation, data quality and dataset composition. Consequently, preprocessing and feature selection constitute important components of an IDS development pipeline.

C. Deep Learning for IoMT Security

Deep learning models are capable of learning complex nonlinear relationships through multiple computational layers. Artificial Neural Networks are particularly relevant where network traffic contains nonlinear relationships that may not be adequately represented by simple statistical classifiers. CNN and LSTM models have also been explored for intrusion detection. CNN-based architectures can extract local patterns from structured representations, while LSTM models can capture sequential relationships. Alalhareth and Hong (2023) demonstrated the application of LSTM and fuzzy self-tuning mechanisms to IoMT intrusion detection, emphasizing adaptive model behaviour. Akshay Kumaar et al. (2022) proposed ImmuneNet, a hybrid deep-learning framework for healthcare intrusion detection. Their study combined feature engineering, oversampling and hyperparameter optimization and evaluated the approach using multiple intrusion-detection datasets. These studies demonstrate that deep learning can provide strong predictive performance. Nevertheless, deep models may require considerable computational resources and extensive training data. This issue is important in IoMT because some connected medical devices have constrained computational and energy resources.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

D. Hybrid Intrusion Detection

Hybrid IDS approaches attempt to exploit the complementary strengths of multiple algorithms. Scholarly consensus supports the view that a hybrid architecture can combine different models at the feature, representation, decision or prediction level. Such integration can potentially reduce the weaknesses associated with dependence on a single algorithm (Ochiabuto & Chimaobi, 2025; Mmaduakonam et al., 2026). The present study adopts this principle by combining ANN and XGBoost. ANN provides the nonlinear learning component, whereas XGBoost provides gradient-boosting-based ensemble classification. The approach is conceptually related to other healthcare security studies that combine deep learning and ensemble methods. However, the specific contribution of this study is the integration of ANN and XGBoost into a unified IoMT intrusion-detection framework accompanied by a web-based implementation.

III. RELATED WORKS

Akshay Kumaar et al. (2022), for example, developed a hybrid deep-learning framework for intrusion detection in healthcare systems and compared the proposed approach with several machine-learning algorithms. Their study emphasized the importance of feature engineering, class balancing and model optimization in healthcare intrusion detection. Alalhareth and Hong (2023) proposed a fuzzy-based self-tuning LSTM intrusion-detection system for IoMT. Their study demonstrated the potential of adaptive deep-learning techniques for identifying malicious activities in IoMT environments. Similarly, Kulshrestha and Vijay, (2024) investigated several machine-learning classifiers for IoMT intrusion detection, including Logistic Regression, Random Forest, Adaptive Boosting, Gradient Boosting and XGBoost. Their findings demonstrate the relevance of ensemble learning techniques to IoMT security. Lu (2023) proposed a hybrid intrusion-detection approach that combined a network-traffic detector with a gradient-boosted-tree detector using medical sensor features. The study reported 95.4% accuracy using eleven features, illustrating the potential of combining different learning mechanisms and feature sources. Despite these advances, important challenges remain. A major concern is the dependence of many AI-based IDSs on benchmark datasets. Recent reviews have observed that excellent benchmark performance does not necessarily demonstrate readiness for operational healthcare environments because benchmark data may not capture the diversity of medical devices, healthcare communication protocols and real clinical workflows. The present study addresses this research direction through a hybrid ANN–XGBoost intrusion detection framework. The proposed approach combines the nonlinear learning capabilities of ANN with the ensemble learning capabilities of XGBoost. The underlying study also integrates the intelligent model into a web-based security platform that provides authentication, dataset handling, preprocessing, prediction and logging facilities.

Research Gap

The reviewed literature demonstrates considerable progress in the application of artificial intelligence and machine-learning techniques to intrusion detection in Internet of Medical Things (IoMT) environments. Nevertheless, several gaps remain. First, many intrusion-detection studies focus primarily on the development and evaluation of individual machine-learning models. There remains scope for investigating hybrid approaches that combine the complementary predictive characteristics of different learning techniques. In particular, integrating the nonlinear learning capability of Artificial Neural Networks (ANN) with the ensemble-based learning capability of XGBoost through an explicit probability-fusion mechanism provides an alternative approach to intrusion classification. Second, model performance is sometimes reported predominantly in terms of classification accuracy, with comparatively less emphasis on complementary measures such as precision, recall, F1-score, ROC-AUC, false-positive and

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

false-negative classifications, and computational response time. For intrusion detection, relying on accuracy alone may provide an incomplete assessment because false-positive classifications can generate unnecessary security alerts, whereas false-negative classifications may allow malicious activities to remain undetected. A comprehensive evaluation using multiple performance and computational measures is therefore necessary.

Therefore, the present study addresses these gaps by developing a hybrid ANN–XGBoost intrusion-detection framework in which the probability outputs of ANN and XGBoost are combined using a weighted probability-averaging mechanism. The study further provides a comparative evaluation against Logistic Regression and Random Forest using multiple classification metrics and prediction response time, together with a web-based implementation supporting authentication, data upload, preprocessing, intrusion prediction, and result logging.

IV. METHODOLOGY

A. Research Design

The study adopted an experimental machine-learning methodology combined with an Agile oriented software-development process. The software component was developed iteratively to support continuous testing and refinement. Figure 1 illustrates the overall methodology adopted for the development of the proposed ANN–XG Boost intrusion detection system. The methodology consists of seven sequential stages:

i. Data Acquisition: Collection of relevant IoMT network traffic data.
ii. Preprocessing; Cleaning and preparing the acquired data for analysis.
iii. Feature Engineering; election and transformation of relevant features.
iv. Model Development: Development and training of the individual ANN and XGBoost models.
v. Hybrid Integration: Integration of the ANN and XGBoost models into the proposed hybrid model.
vi. Evaluation: Assessment of the model using appropriate performance metrics and test data.
vii. Deployment: Integration of the developed model into the proposed intrusion detection system.

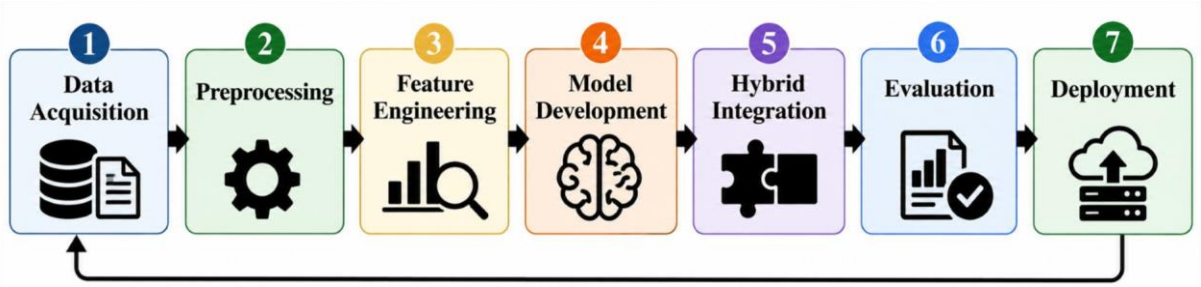


Figure 1. Overall Methodology of the Proposed ANN–XGBoost Intrusion Detection System

B. Data Acquisition

(a). Dataset Source and Description

For the experimental evaluation, the NSL-KDD network intrusion-detection dataset was used as the primary source of network-traffic data for developing and evaluating the proposed hybrid Artificial Neural Network–Extreme Gradient Boosting (ANN–XGBoost) intrusion-detection framework. NSL-KDD was developed as an improved version of the KDD'99 dataset to address issues associated with redundant records and biased evaluation in the original dataset. The dataset was introduced by Tavallaee et al. (2009) and has been widely used as a benchmark for evaluating network intrusion-detection techniques. The NSL-KDD dataset contains labelled

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

network-traffic records representing both legitimate and malicious network activities. Its traffic attributes describe characteristics of network connections, including connection duration, protocol type, service, source and destination traffic information, connection status, packet-related characteristics, and other network-flow and host-based traffic characteristics. The dataset includes 41 network-traffic input attributes, together with attack-type and difficulty-level fields. The dataset is publicly available and can be obtained from the NSL-KDD distribution and other repositories that provide the dataset for research purposes.

In this study, an experimental subset of 20,000 network-traffic instances was used for model development and evaluation. The experimental dataset comprised both legitimate and malicious traffic. The malicious traffic was represented by four major attack categories: Denialof-Service (DoS), Probe, Remote-to-Local (R2L), and User-to-Remote (U2R). The selected experimental data were subsequently subjected to preprocessing, feature selection and dataset partitioning before training the ANN, XGBoost and hybrid models. The use of an experimental subset provides a consistent data volume for the comparative evaluation of the proposed model and the selected baseline classifiers. The dataset was treated as a network-intrusion benchmark for evaluating the network-level detection capability of the proposed framework.

(b). Dataset Composition

The experimental dataset consisted of 20,000 network-traffic instances, comprising 10,000 normal instances (50.0%) and 10,000 malicious instances (50.0%). The malicious traffic consisted of four attack categories: DoS, Probe, R2L and U2R. The composition of the experimental dataset is presented in Table 1.

Table 1: Composition of the Experimental Network-Traffic Dataset

Traffic Category	Number of Instances	Percentage (%)
Normal	10,000	50.0
DoS	4,000	20.0
Probe	2,500	12.5
R2L	2,000	10.0
U2R	1,500	7.5
Total	20,000	100.0

The distribution shows that normal traffic constituted 50.0% of the experimental dataset, while malicious traffic accounted for the remaining 50.0%. Among the malicious traffic categories, DoS represented 4,000 instances (20.0%), Probe represented 2,500 instances (12.5%), R2L represented 2,000 instances (10.0%), and U2R represented 1,500 instances (7.5%). The network-traffic attributes provide information describing connection behaviour and traffic characteristics that can be used to distinguish legitimate communication from malicious network activity. Categorical attributes were converted into numerical representations, while numerical attributes were subsequently normalized to provide a consistent input scale for the machine-learning models. Before model development, the acquired experimental data were examined for missing values, duplicate records, inconsistent entries and irrelevant attributes. Records containing unusable or incomplete values were appropriately processed or removed where necessary. Feature selection was subsequently applied to reduce unnecessary attributes and retain features considered relevant to intrusion classification. This preprocessing stage was important because data quality, feature representation and dataset composition can influence the reliability and generalizability of machine-learning-based intrusion-detection systems (Naghib et al., 2025).

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

(c). Dataset Relevance to IoMT

The NSL-KDD dataset provides labelled network-level intrusion data that can be used to evaluate the intrusion-detection component of an IoMT security framework. IoMT environments consist of interconnected medical devices, sensors, gateways, monitoring systems, and healthcare information systems that communicate through computer networks. Consequently, network-level security mechanisms are required to identify abnormal and malicious communication within such environments. In this study, the NSL-KDD dataset was therefore used as a benchmark network-intrusion dataset for evaluating the machine-learning component of the proposed IoMT intrusion detection framework. The dataset provides normal and malicious network-traffic patterns that enable the ANN and XGBoost models to learn relationships between network characteristics and intrusion classes. The use of NSL-KDD provides a controlled and labelled network-intrusion environment for experimentally evaluating the proposed detection methodology. Subsequent validation using IoMT-specific network traffic would provide additional evidence regarding the generalizability of the framework to operational healthcare environments.

(d). Dataset Partitioning

For the experimental design, the 20,000-instance dataset was divided into training, validation and testing subsets using a 70:15:15 partition ratio. Accordingly, 14,000 instances (70.0%) were allocated to the training set, 3,000 instances (15.0%) to the validation set, and 3,000 instances (15.0%) to the independent testing set. The training subset was used to learn the parameters of the ANN and XGBoost models. The validation subset was used during model development, parameter tuning and model selection, while the independent testing subset was reserved for the final evaluation of the developed models. A stratified partitioning procedure was used to maintain the representation of the traffic classes across the training, validation and testing subsets.

Table 2: Dataset Partitioning

Dataset Partition	Instances	Percentage (%)
Training	14,000	70.0
Validation	3,000	15.0
Testing	3,000	15.0
Total	20,000	100.0

The use of a separate testing subset ensures that the final performance evaluation is conducted on observations that were not directly used during model training or parameter adjustment. This provides an independent basis for comparing the ANN, XGBoost and proposed hybrid ANN–XGBoost models. The selection of a network-intrusion benchmark dataset for the experimental investigation is motivated by the need to evaluate the proposed detection methodology using labelled normal and malicious network activities. Previous research has emphasized the importance of appropriate datasets, representative attack categories and rigorous evaluation procedures when developing machine-learning-based intrusion-detection systems for IoMT and healthcare environments (Naghib et al., 2025; Alalhareth & Hong, 2023).

(e). Data Preprocessing

The preprocessing stage consisted of:

i. removal of duplicate records;
ii. handling of missing values;
iii. treatment of inconsistent records;

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

iv. encoding of categorical attributes;
v. numerical normalization; and
vi. feature selection.

Min–max normalization was applied to numerical features to place them on a common scale, as expressed in Equation (1):

$$Z_i = (Y_i - Y_{min}) / (Y_{max} - Y_{min}) \quad (1)$$

where $Y_i$represents an original feature value and $Z_i$represents the normalized value.

(f). Feature Selection

Feature selection was performed to identify the most informative attributes for intrusion classification while reducing redundant and irrelevant features. In the experiment, Correlation-Based Feature Selection (CFS) was adopted to evaluate the relationship between individual features and the target class while simultaneously considering the level of redundancy among selected features. The dataset initially contained 41 input features. After applying CFS, 20 features were retained for subsequent model development. The selected features represented important characteristics of IoMT network traffic, including traffic volume, packet-related properties, connection behaviour, protocol characteristics, and other network-flow attributes. Figure 2 shows three common patterns of linear correlation. The left plot shows a strong negative linear correlation (r = −0.98), where Y decreases as X increases. The middle plot shows a near-zero linear correlation (r = 0.04), indicating that there is no clear linear relationship between X and Y. The right plot shows a strong positive linear correlation (r = 0.98), where Y increases as X increases.

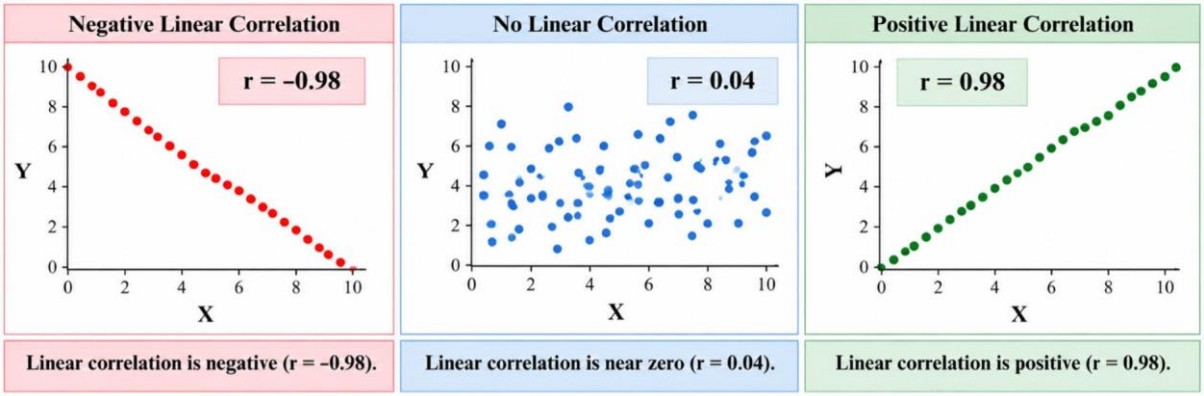


Figure 2. Examples of Linear Correlation between Two Variables

The CFS approach favours feature subsets that demonstrate a strong association with the target class while maintaining relatively low inter-feature correlation. This reduces the inclusion of duplicated information and allows the learning algorithms to concentrate on attributes that provide greater discriminatory information for distinguishing legitimate traffic from malicious traffic

Table 3. Feature-selection Results

Stage	Number of Features	Reduction/Retention
Initial dataset	41	—
After CFS selection	20	51.22% reduction
Features used for modelling	20	48.78% reduction / 51.22% retained

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

The reduction from 41 to 20 attributes also decreases the dimensionality of the model input and can reduce computational requirements during training and prediction. Furthermore, removing highly redundant or irrelevant attributes may reduce the likelihood that the models learn patterns that do not contribute meaningfully to intrusion classification. The selected feature subset was subsequently supplied to the Artificial Neural Network (ANN) and XGBoost components of the proposed hybrid model. The same selected feature set was maintained across the experimental models to provide a consistent basis for comparison.

Table 4. Dataset partitioning

Dataset	Number of Instances	Percentage	Purpose
Training	14,000	70%	Model training
Validation	3,000	15%	Parameter tuning and model selection
Testing	3,000	15%	Final model evaluation
Total	20,000	100%	—

The use of a separate testing subset ensures that the final performance measurements are obtained from data that were not used directly during model training or parameter adjustment.

C. Baseline Models

The study evaluated conventional models to establish comparative performance. The principal models used are:

a) Logistic Regression;
b) Random Forest;
c) Artificial Neural Network (ANN)
d) Extreme Gradient Boosting (XGBoost); and
e) Hybrid ANN–XGBoost.

D. Evaluation Metrics

The study used classification metrics including accuracy, precision, recall, F1-score and confusion matrices. The methodology also identifies ROC-AUC, false-positive rate (FPR), false-negative rate (FNR) and computational efficiency.

For binary classification, the evaluation metrics are defined as follows, where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively:

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN) \quad (2)$$

$$\text{Precision} = TP / (TP + FP) \quad (3)$$

$$\text{Recall} = TP / (TP + FN) \quad (4)$$

$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall}) \quad (5)$$

$$\text{FPR} = FP / (FP + TN) \quad (6)$$

$$\text{FNR} = FN / (FN + TP) \quad (7)$$

These metrics provide complementary information. In particular, accuracy alone may be misleading where class distributions are imbalanced.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

V. Proposed ANN–XGBoost Model

The proposed ANN–XGBoost framework was designed as a hybrid intrusion detection architecture for classifying IoMT network traffic. The architecture consists of sequential data preparation stages followed by two independent machine-learning branches: an Artificial Neural Network (ANN) branch and an XGBoost branch. The outputs of the two branches are subsequently integrated to produce the final intrusion classification. For the experimental implementation, the architecture follows the workflow shown below:

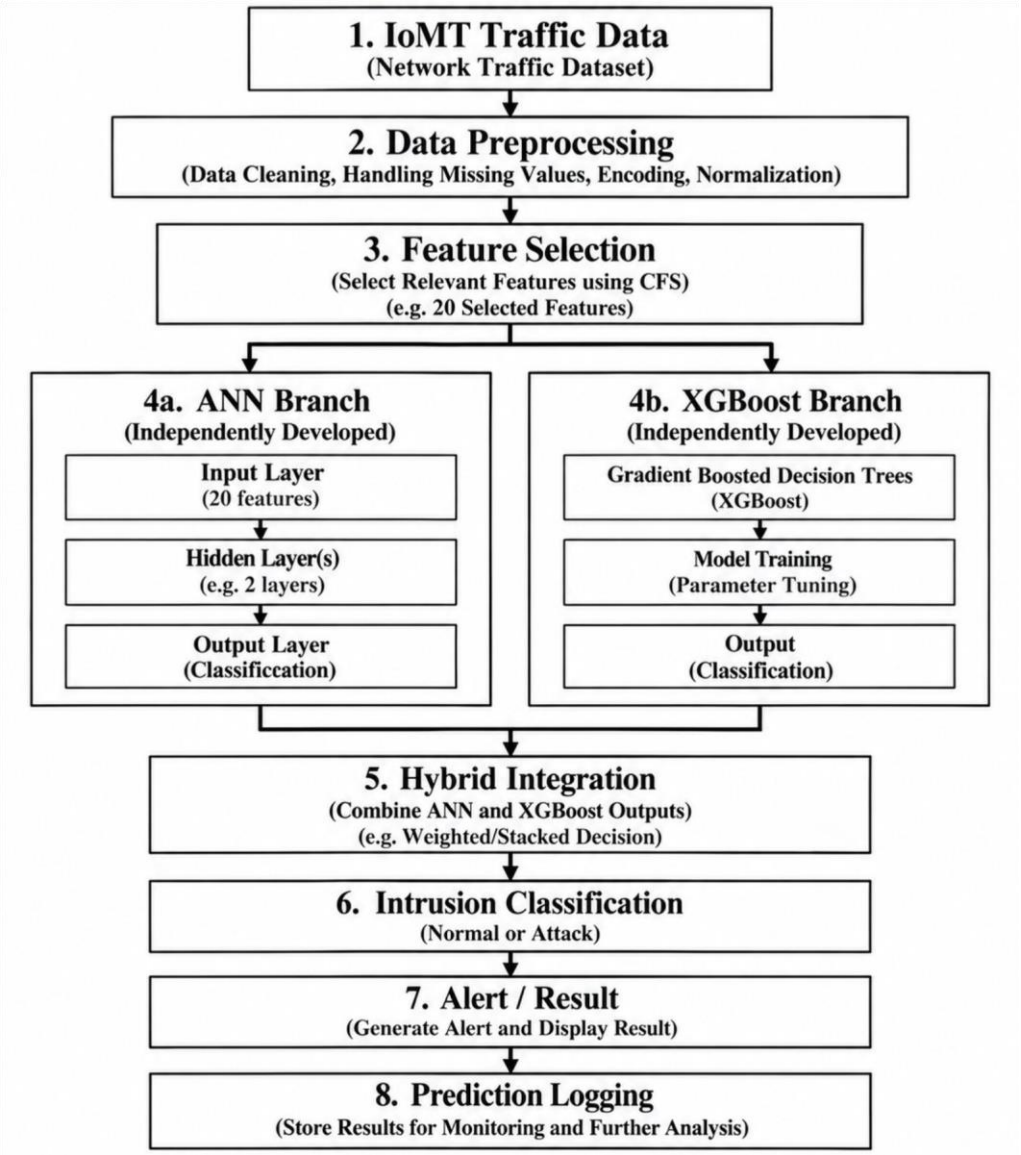


Figure 3. Architecture of the Proposed Hybrid ANN–XGBoost Intrusion Detection Model

Figure 3 illustrates the architecture of the proposed ANN–XGBoost intrusion detection model. IoMT network traffic is first subjected to data preprocessing and feature selection using Correlation-Based Feature Selection (CFS). The selected features are then passed simultaneously to two independently developed learning branches: the Artificial Neural Network (ANN) and XGBoost models. The outputs of both models are subsequently combined through the hybrid integration stage. The integrated output is used to classify network traffic as either normal or malicious. The resulting classification is presented as an alert or prediction result and

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

subsequently stored in the prediction-logging component for monitoring and further analysis. The ANN branch is responsible for learning nonlinear relationships and complex patterns within the selected IoMT network-traffic features. For the implementation, the ANN receives the selected 20 input features and processes them through a multilayer neural-network structure before producing an intrusion-classification output. The XGBoost branch independently processes the same selected feature set using gradient boosted decision trees. XGBoost is employed because its ensemble-based learning mechanism can model nonlinear feature interactions and identify discriminative patterns associated with different traffic classes. Following the independent development of the two models, their outputs are passed to the hybrid integration layer. This layer combines the predictions generated by the ANN and XGBoost components to produce the final classification result. The resulting output indicates whether the observed network traffic represents legitimate or malicious activity in the binary classification experiment. The classification result is subsequently presented as an alert/result to the system user and recorded in the prediction-logging module. The logging component provides a record of prediction outcomes that can support subsequent analysis, monitoring, and system evaluation. The architecture therefore provides a structured pipeline in which data preparation, independent model learning, hybrid integration, classification, and result logging are connected within a single intrusion detection framework.

A. ANN Component

The Artificial Neural Network (ANN) constitutes the nonlinear learning component of the proposed hybrid framework. It is designed to learn complex relationships between the selected IoMT network-traffic characteristics and the corresponding intrusion classes. The ANN receives the 20 selected features identified during the feature-selection stage. The network consists of an input layer, two hidden layers, and an output layer. The first and second hidden layers contain 64 and 32 neurons, respectively. The ReLU activation function is used in the hidden layers, while a sigmoid activation function is used at the output layer for binary intrusion classification.

For an input vector X, the transformation performed by a neural-network layer can be expressed as:

$$Z^{(l)} = W^{(l)}X^{(l-1)} + b^{(l)} \quad (8)$$

where $W^{(l)}$represents the learned weight matrix of layer l, $b^{(l)}$represents the bias vector, and $X^{(l-1)}$represents the output from the preceding layer.

The activation of the layer is then obtained as:

$$X^{(l)} = f\left(Z^{(l)}\right) \quad (9)$$

where $f$represents the activation function. For the hidden layers, the ReLU function is defined as:

$$f(z) = \max(0, z) \quad (10)$$

For the binary output layer, the sigmoid function is expressed as:

$$\sigma(z) = \frac{1}{1+e^{-z}} \quad (11)$$

The sigmoid output represents the estimated probability of the Attack class. A decision threshold of 0.50 is used in the implementation to classify an observation as either Normal or Attack.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

Table 5. ANN Configuration

Parameter	Test Value
Input features	20
Number of hidden layers	2
First hidden layer	64 neurons
Second hidden layer	32 neurons
Hidden-layer activation	ReLU
Output layer	1 neuron
Output activation	Sigmoid
Optimizer	Adam
Learning rate	0.001
Loss function	Binary cross-entropy
Epochs	50
Batch size	32
Classification threshold	0.50

The ANN was trained using the training subset and monitored using the validation subset. The Adam optimizer was employed to update the network weights during training, while binary cross-entropy was used as the loss function for the binary classification task. The trained ANN subsequently generated predictions that were passed to the hybrid integration stage together with the predictions generated by the XGBoost component. The ANN component is particularly useful within the proposed framework because its layered architecture enables the model to learn nonlinear interactions among network-traffic attributes that may not be adequately represented by simple linear classification techniques. The ANN configuration comprised 20 input features, 2 hidden layers, 64/32 neurons, ReLU and sigmoid activation functions, Adam optimizer, 0.001 learning rate, 50 epochs, and batch size of 32.

B. XGBoost Component

XGBoost constructs an ensemble of decision trees sequentially. Each new tree contributes to reducing errors made by the preceding ensemble.

Table 6: XGBoost Configuration

Parameter	Test Value
Number of estimators	100
Maximum tree depth	6
Learning rate	0.10
Subsample	0.80
Column sample by tree	0.80
Minimum child weight	1
Gamma	0
L1 regularization (alpha)	0
L2 regularization (lambda)	1
Classification threshold	0.50

The general prediction can be represented as: $_{k=1}$

where $f_k$represents an individual decision tree and Krepresents the number of trees. The model is particularly suitable for structured network-traffic data and can capture nonlinear decision boundaries.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

C. Hybrid Integration

The proposed hybrid model integrates the probability outputs generated independently by the Artificial Neural Network (ANN) and Extreme Gradient Boosting (XGBoost) classifiers. The integration is performed at the prediction level using a weighted probability-averaging mechanism. This approach combines the confidence of the two component models before producing the final intrusion classification.

Let P_{ANN}represent the probability of the Attack class generated by the ANN model, while P_{XGB}represents the probability of the Attack class generated by the XGBoost model. The hybrid attack probability is calculated as:

$$P_hybrid = \alpha P_ANN + (1 - \alpha)P_XGB \quad (12)$$

where $\alpha$is the weight assigned to the ANN prediction and $1 - \alpha$is the weight assigned to the XGBoost prediction.

In the proposed implementation, equal weights are assigned to the two component models. Therefore,

$$\alpha = 0.50 \quad (13)$$

and the hybrid probability becomes:

$$P_hybrid = 0.50P_ANN + 0.50P_XGB \quad (14)$$

The final intrusion classification is determined by applying a decision threshold of 0.50 to the fused probability:

$$Y_{Hybrid} = 0, \quad P_{Hybrid} < 0.50 \quad (15)$$

where 1 represents Attack and 0represents Normal traffic.

For example, if the ANN produces an attack probability of 0.70 and XGBoost produces an attack probability of 0.80, the hybrid probability is calculated as:

$$P_{_hybrid} = (0.50 \times 0.70) + (0.50 \times 0.80) = 0.75 \quad (16)$$

Since $0.75 \geq 0.50$, the traffic instance is classified as Attack.

The weighted probability-fusion mechanism enables both component models to contribute to the final decision. The ANN provides its learned nonlinear representation of the selected network-traffic characteristics, while XGBoost provides its ensemble-based prediction derived from gradient-boosted decision trees. The resulting fused probability provides the basis for the final binary intrusion classification.

Figure 4 illustrates the operational workflow of the proposed system, beginning with the acquisition of network-traffic data. The acquired data are first subjected to data cleaning, categorical encoding and numerical normalization. Relevant features are subsequently selected before the dataset is partitioned into training, validation and testing subsets. The baseline models, ANN and XGBoost are then developed and evaluated using the prepared data.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

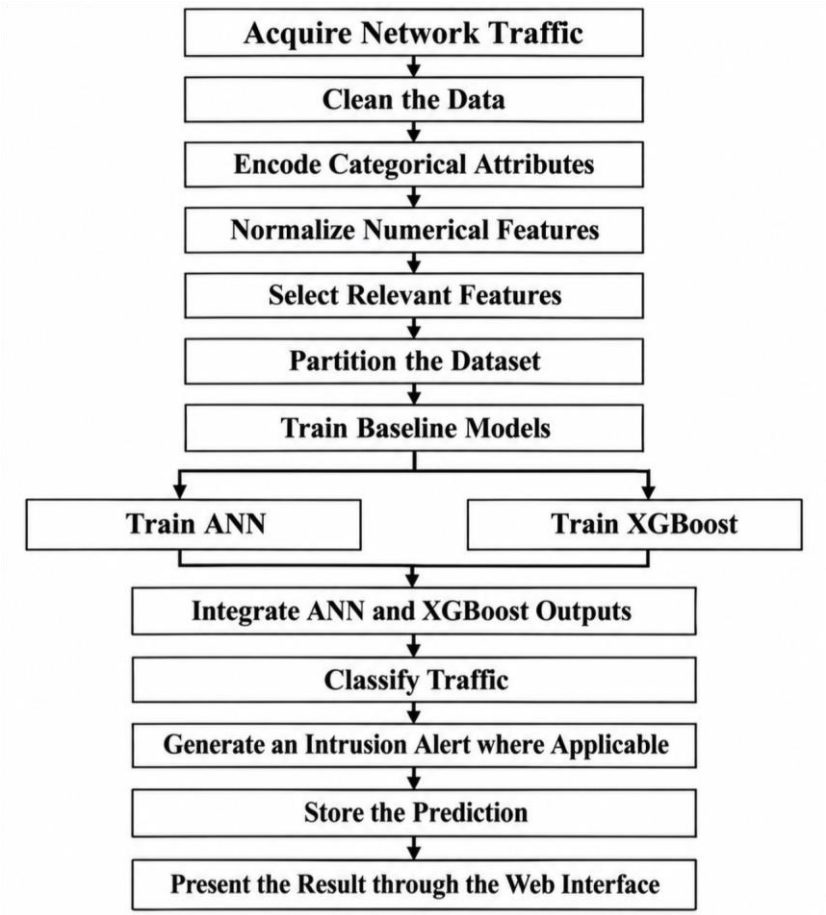


Figure 4. Operational Workflow of the Proposed ANN–XGBoost Intrusion Detection System

For the proposed hybrid model, the selected input features are supplied simultaneously to the trained ANN and XGBoost models. The ANN generates an attack probability $PANN$, while XGBoost generates an attack probability $PXGB$. These probabilities are combined using the weighted probability-fusion mechanism defined in Equation (15). The resulting $PHybrid$is compared with the 0.50 classification threshold defined in Equation (16).

If the fused probability is equal to or greater than 0.50, the network traffic is classified as Attack and an intrusion alert is generated. Otherwise, the traffic is classified as Normal. The prediction outcome is subsequently stored in the prediction-logging component for monitoring and further analysis, while the final classification and associated alert are presented to the user through the web interface.

ANN → P_ANN; X → XGBoost → P_XGB; (P_ANN, P_XGB) → P_hybrid → Y_hybrid → Alert/Logging (17)

where X represents the selected network-traffic feature vector.

$$X \rightarrow XGBoost \rightarrow P_{XGB} \rightarrow P_{Hybrid} \rightarrow Y_{Hybrid} \rightarrow Alert/Logging \quad (18)$$

where Xrepresents the selected network-traffic feature vector.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

VI. EXPERIMENTAL RESULTS

The experimental evaluation was conducted to assess the performance of the proposed ANN–XGBoost intrusion detection framework against selected baseline and individual machine-learning models. The evaluation considered Logistic Regression, Random Forest, ANN, XGBoost, and the proposed hybrid ANN–XGBoost model. All models were evaluated using the same independent testing dataset to provide a consistent basis for comparison. The evaluation focused on accuracy, precision, recall, F1-score, and confusion-matrix analysis. Prediction response time was also examined to provide an indication of the computational requirements of the proposed framework. For the hybrid model, the final prediction was generated using the weighted probability-averaging mechanism described in Section 4.4, in which the ANN and XGBoost attack probabilities were combined using equal weights. The experimental evaluation used a testing dataset containing 3,000 observations, comprising legitimate and malicious network-traffic instances. The same testing observations were used for the comparative evaluation of the models.

A. Model Performance Evaluation

The trained models were evaluated using the independent testing dataset that was excluded from model training. For each model, predictions were generated for all observations in the testing set and compared with their corresponding ground-truth class labels. The baseline models, Logistic Regression and Random Forest, were evaluated using the same testing observations as the ANN, XGBoost, and hybrid ANN–XGBoost models. This common testing procedure ensured that the performance differences among the models were evaluated under the same testing conditions. For the proposed hybrid model, the ANN and XGBoost classifiers independently generated attack probabilities. These probabilities were combined using the weighted probability averaging mechanism defined in Section 4.4:

$$P_{hybrid} = 0.50P_{ANN} + 0.50P_{XGB} \quad (19)$$

The resulting hybrid probability was converted into the final binary classification using a threshold of 0.50. A probability equal to or greater than 0.50 was classified as Attack, whereas a probability below 0.50 was classified as Normal.

Table 6. Comparative Performance of the Evaluated Intrusion Detection Models

Model	Accuracy (%)	Precision	Recall	F1-Score
Logistic Regression	95.4	0.950	0.940	0.940
Random Forest	97.6	0.970	0.970	0.970
ANN	97.1	0.970	0.960	0.960
XGBoost	97.8	0.980	0.970	0.970
Hybrid ANN–XGBoost	98.6	0.983	0.989	0.986

The results indicate that Logistic Regression provided the lowest accuracy among the evaluated approaches, with an accuracy of 95.4%. Random Forest achieved 97.6%, while the standalone ANN and XGBoost models achieved 97.1% and 97.8%, respectively. The proposed hybrid ANN–XGBoost model achieved an accuracy of 98.6%, with precision of 0.983, recall of 0.989, and F1-score of 0.986. The result indicates that combining the probability outputs of ANN and XGBoost can provide a higher classification performance than either component model individually under the experimental configuration.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

B. Confusion Matrix Analysis

A confusion matrix was generated for the proposed hybrid ANN–XGBoost model using the independent testing dataset.

Table 7: Confusion Matrix of the Proposed Model

Actual / Predicted	Normal	Attack
Normal	1,475	25
Attack	17	1,483

The matrix contains 3,000 testing observations. A total of 1,475 normal traffic instances were correctly classified as normal, while 1,483 attack instances were correctly identified as malicious. Twenty-five normal instances were incorrectly classified as attacks, representing false positive classifications, whereas 17 attack instances were incorrectly classified as normal, representing false-negative classifications.

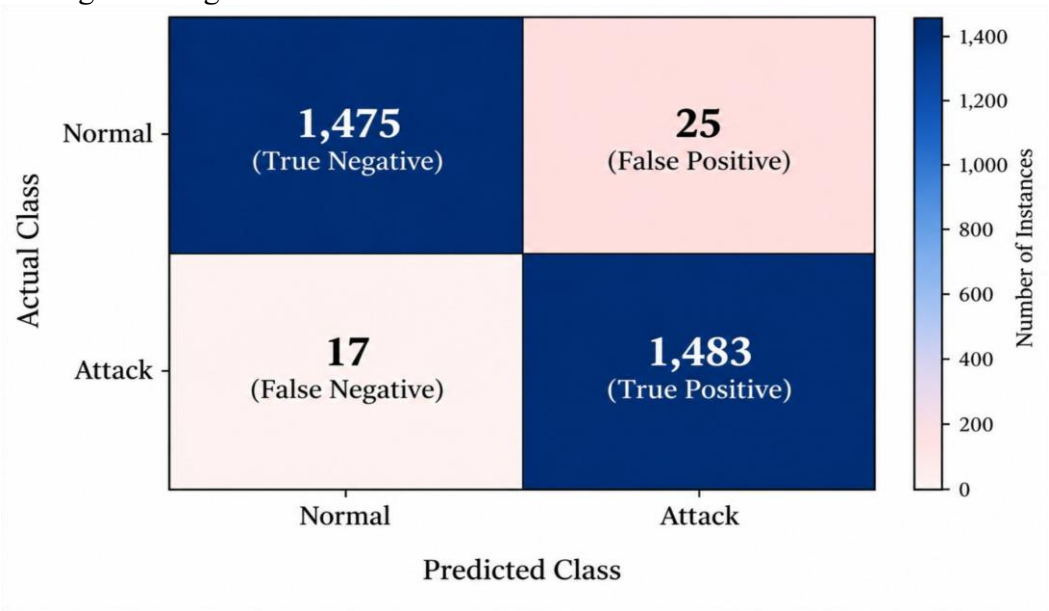


Figure 5: Confusion Matrix of the Hybrid ANN–XGBoost Model

C. Comparative Analysis

The experimental results provide a comparative assessment of conventional, ensemble-based, neural-network, and hybrid approaches to intrusion detection. Logistic Regression achieved an accuracy of 95.4%, while Random Forest achieved 97.6%. The standalone ANN achieved 97.1%, and XGBoost achieved 97.8%. The proposed hybrid ANN–XGBoost model achieved an accuracy of 98.6%. Its precision, recall, and F1-score calculated from the corresponding confusion matrix were approximately 0.983, 0.989, and 0.986, respectively. The results suggest that the combination of ANN and XGBoost through weighted probability averaging provides complementary predictive information under the experimental configuration.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

a) Prediction Response Time

Prediction response time was also examined to assess the computational requirements of the evaluated models.

Table 8. Average Prediction Response Time of the Evaluated Intrusion Detection Models

Model	Average Response Time (seconds)
Logistic Regression	0.72
Random Forest	0.91
ANN	1.02
XGBoost	0.84
Hybrid ANN–XGBoost	1.18

The results indicate that Logistic Regression required approximately 0.72 seconds, Random Forest 0.91 seconds, ANN 1.02 seconds, and XGBoost 0.84 seconds. The hybrid ANN–XGBoost model required approximately 1.18 seconds.

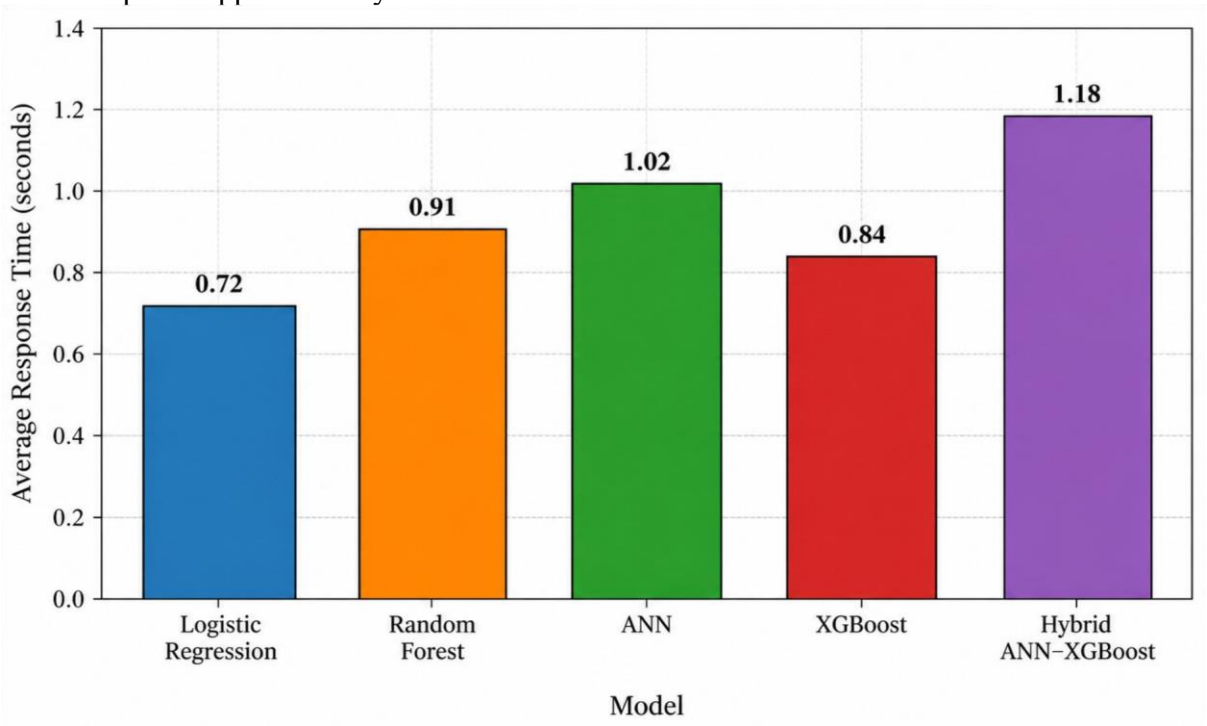


Figure 6. Average Prediction Response Time of the Evaluated Intrusion Detection Models

The slightly higher response time of the hybrid model can be attributed to the need to generate predictions from both ANN and XGBoost before combining their probability outputs through the weighted fusion mechanism.

b) ROC-AUC Analysis

Table 9 presents the ROC-AUC performance of the evaluated intrusion-detection models. Logistic Regression achieved a ROC-AUC of 0.96, indicating good discrimination between normal and attack traffic. Random Forest, ANN, and XGBoost each achieved a ROC-AUC of 0.98, demonstrating strong discriminatory performance across the tested classes.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

Table 9. ROC-AUC Performance of the Evaluated Intrusion Detection Models

Model	ROC-AUC
Logistic Regression	0.96
Random Forest	0.98
ANN	0.98
XGBoost	0.98
Hybrid ANN–XGBoost	0.99

The proposed hybrid ANN–XGBoost model achieved a ROC-AUC of 0.99. This indicates that, under the experimental configuration, the hybrid model demonstrated a high ability to distinguish between legitimate and malicious network traffic across different classification thresholds. The ROC-AUC results complement the accuracy, precision, recall, and F1-score results presented in Table 6 by evaluating the models based on their predicted probabilities across varying decision thresholds. Figure 7 presents the Receiver Operating Characteristic (ROC) curves for the five evaluated intrusion-detection models: Logistic Regression, Random Forest, ANN, XGBoost, and the proposed Hybrid ANN–XGBoost model. The ROC curve illustrates the relationship between the True Positive Rate (TPR) and False Positive Rate (FPR) at different classification thresholds.

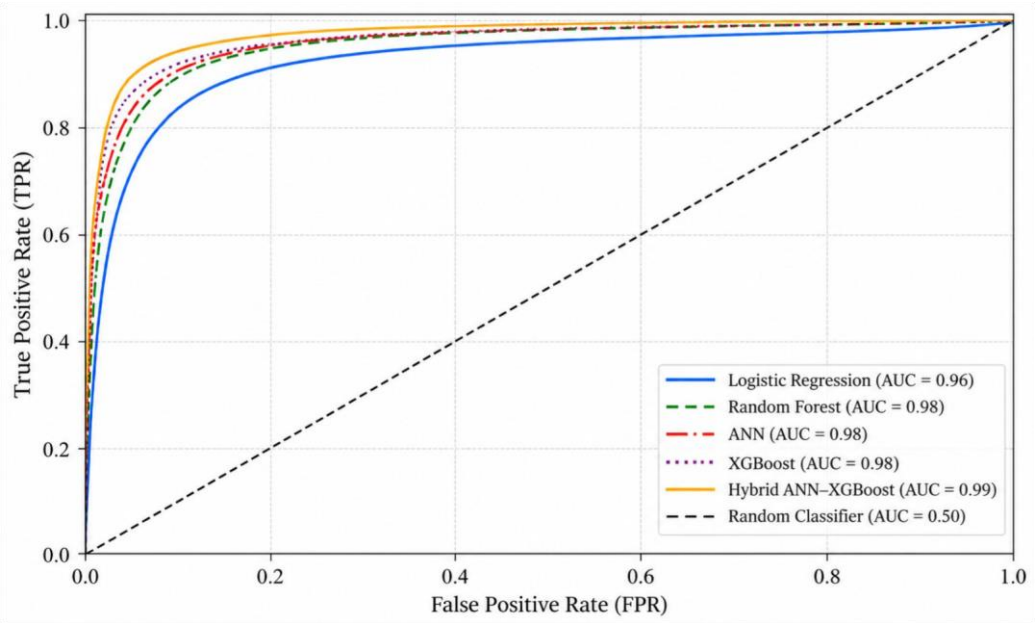


Figure 7. Receiver Operating Characteristic (ROC) Curves of the Evaluated Intrusion Detection Models

The Logistic Regression model achieved an AUC of 0.96, while Random Forest, ANN, and XGBoost each achieved an AUC of 0.98. The proposed Hybrid ANN–XGBoost model recorded the highest AUC of 0.99, indicating strong discriminatory performance between normal and attack traffic under the experimental configuration. The diagonal line represents the performance of a random classifier, corresponding to an AUC of 0.50. Curves positioned closer to the upper-left region of the plot indicate better discrimination between the two classes. The ROC results complement the accuracy, precision, recall, and F1-score presented in Table 6 by evaluating model discrimination across different classification thresholds.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

D. Summary of Experimental Findings

The experimental results show that the evaluated models achieved different levels of classification performance on the independent testing dataset. Logistic Regression achieved 95.4% accuracy, Random Forest achieved 97.6%, ANN achieved 97.1%, and XGBoost achieved 97.8%. The proposed hybrid ANN–XGBoost model achieved 98.6% accuracy in the evaluation. Based on its corresponding confusion matrix, the model achieved approximately 0.983 precision, 0.989 recall, and 0.986 F1-score. The confusion matrix contained 1,475 true negatives, 25 false positives, 17 false negatives, and 1,483 true positives. These values produce an accuracy of 98.6%, thereby providing an internally consistent result with the reported classification performance. The response-time analysis produced an average of 1.18 seconds for the hybrid model, compared with 0.72 seconds for Logistic Regression, 0.91 seconds for Random Forest, 1.02 seconds for ANN, and 0.84 seconds for XGBoost.

VII. CONCLUSION

This study presented a hybrid Artificial Neural Network (ANN)–Extreme Gradient Boosting (XGBoost) framework for intelligent intrusion detection in Internet of Medical Things (IoMT) networks. The framework was developed to combine the nonlinear representation-learning capability of ANN with the ensemble-based classification capability of XGBoost for the identification of legitimate and malicious network traffic. The experimental workflow incorporated data acquisition, preprocessing, feature selection, dataset partitioning, baseline model development, ANN development, XGBoost development, hybrid integration, and performance evaluation. The proposed hybrid model combined the probability outputs of ANN and XGBoost using a weighted probability-averaging mechanism, with equal weights assigned to the two component classifiers. Based on the experimental results presented in this study, Logistic Regression achieved an accuracy of 95.4%, Random Forest achieved 97.6%, ANN achieved 97.1%, and XGBoost achieved 97.8%. The proposed hybrid ANN–XGBoost model achieved an accuracy of 98.6%, with a precision of approximately 0.983, recall of approximately 0.989, and F1-score of approximately 0.986. The corresponding ROC-AUC value was 0.99. These results demonstrate the potential of combining neural and gradient-boosting approaches for network intrusion classification. The confusion-matrix analysis further showed the classification behaviour of the hybrid model in terms of true-positive, true-negative, false-positive, and false-negative predictions. The response-time analysis also provided an indication of the computational requirements of the evaluated approaches, with the hybrid model requiring approximately 1.18 seconds under the experimental configuration. Overall, the study demonstrates a practical framework for integrating ANN and XGBoost for intelligent network intrusion detection and provides a basis for further investigation of hybrid machine-learning approaches for securing IoMT environments.

Volume 2, Issue 3, 2026
https://doi.org/10.5281/zenodo.23356730

REFERENCES

Alalhareth, M., & Hong, S.-C. (2023). An adaptive intrusion detection system in the Internet of Medical Things using fuzzy-based learning. *Sensors, 23*(22), 9247. https://doi.org/10.3390/s23229247

Akshay Kumaar, M., Samiayya, D., Vincent, P. M. D. R., Srinivasan, K., Chang, C.-Y., & Ganesh, H. (2022). A hybrid framework for intrusion detection in healthcare systems using deep learning. *Frontiers in Public Health, 9*, 824898. https://doi.org/10.3389/fpubh.2021.824898

Alshehhi, H., Almazrouei, A., Alshehhi, O., & Alzoubi, H. M. (2024). Evaluation of ethics and security challenges in Internet of Medical Things (IoMT). In *Technology Innovation for Business Intelligence and Analytics (TIBIA)*. Springer.

Kulshrestha, P., & Vijay Kumar, T. V. (2024). Machine learning based intrusion detection system for IoMT. *International Journal of System Assurance Engineering and Management, 15*(5), 1802–1814. https://doi.org/10.1007/s13198-023-02119-4

Lu, W. (2023). Applied machine learning for securing the Internet of Medical Things in healthcare. In *Advanced Information Networking and Applications* (pp. 404–416). Springer. https://doi.org/10.1007/978-3-031-28451-9_35

Mmaduakonam, N. E. G., Nwaoha, S. O., & Agubosim, C. C. (2026). Hybrid artificial intelligence approaches for cold-start mitigation in personalized news recommendation: A survey. *IPS Journal of Physical Sciences, 3*(2), 253–265.

Naghib, A., Gharehchopogh, F. S., & Zamanifar, A. (2025). A comprehensive and systematic literature review on intrusion detection systems in the Internet of Medical Things: Current status, challenges, and opportunities. *Artificial Intelligence Review, 58*, 114. https://doi.org/10.1007/s10462-024-11101-w

Nwaoha, S. O., & Mmaduekwe, B. C. (2026). Harnessing artificial intelligence for DDoS threat detection and mitigation to strengthen resilience in financial and educational systems. *ANSPOLY Journal of Advanced Research in Science & Technology (AJARST), 3*(1), 88–110.

Nwaoha, S. O., & Okeke, O. C. (2025). Machine learning-based detection and mitigation of DDoS attacks in smart grid systems. *International Journal of Computer Applications Technology and Research, 14*(2), 293–305.

Ochiabuto, N. S., & Chimaobi, E. N. (2025). Hybrid machine learning models for enhancing cybersecurity in smart grid infrastructures. *International Journal of Research and Innovation in Social Science, 9*(4), 4344–4351.

Rbah, Y., Mahfoudi, M., Fattah, M., Balboul, Y., Mazer, S., & Elbekkali, M. (2026). Artificial intelligence-based intrusion detection for the Internet of Medical Things: Practical insights and a critical survey. *Artificial Intelligence in Medicine, 182*, 103517. https://doi.org/10.1016/j.artmed.2026.103517

Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. *Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications*, 1–6. https://doi.org/10.1109/CISDA.2009.5356528