

ENSEMBLE MACHINE LEARNING MODELS FOR INTRUSION DETECTION USING THE NSL-KDD DATASET

Deep Shaileshbhai Patel (2504308)
ARTIFICIAL INTELLIGENCE (ACCA7026)

Table of Contents

Abstract	2
1. Introduction.....	3
2. Problem Statement.....	4
3. Related Work	5
4. Methodology	7
4.1 Acquisition of Dataset	7
4.2 Data Preprocessing	7
4.3 Model Development.....	8
4.4 Model Assessment	8
4.5 Analysis of Feature Importance.....	8
4.6 Interpretation and Creation of Insight.....	8
5. Results and Discussion	9
5.1 Classification Performance	9
5.2 Interpretation of the Confusion Matrix	9
5.3 Efficiency of Computation	11
5.4 Feature Importance Analysis.....	11
5.5 Overview of the Discussion.....	11
6. Conclusion and Recommendations	14
6.1 Recommendations.....	14
6.2 Limitations	15
7. References.....	16

Abstract

One of the primary issues with cybersecurity is still network intrusion detection [1]. To ensure that multi-class intrusion detection, this study preprocesses the NSL-KDD dataset and applies two ensemble classifiers: Random Forest (bagging) and XGBoost (gradient boosting) [2], [4], [5]. We evaluate computational efficiency and robustness, analyse feature importances, assess standard classification metrics (accuracy, precision, recall, F1-scores, ROC-AUC where applicable), tune hyperparameters, and interpret the results to offer suggestions for IDS design. While Random Forest produces competitive results with easier customisation and frequently better interpretability for specific feature subsets, XGBoost exhibits stronger discriminative performance in structured traffic data [5], [14]. Features like service, src_bytes, logged_in and flag were found to be important factors in intrusion detection through feature importance analysis. The results imply that feature relevance insights can direct the creation of intelligent and adaptive network security systems, and ensemble models can greatly enhance IDS performance [6], [14]. This architecture provides a strong and dependable solution for contemporary network intrusion detection by skillfully fusing adaptive ensemble learning with class-balancing techniques [2], [11], [14].

The results show that Gradient Boosting outperformed Random Forest (0.9192) with an accuracy of 0.9308, especially in minority attack detection. These findings encourage more research into adaptive and hybrid intrusion detection models and show that ensemble approaches are appropriate for intelligent IDS design [5], [11], [14].

1. Introduction

Our increasing reliance on digital technologies has resulted in more cyber threats attacking networks, data, and vital infrastructure [1], [18]. Intrusion detection systems are essential for monitoring network settings by identifying unusual and malicious activity. Traditional IDS models fail to detect unexpected or zero-day attacks because of their dependence on static rules and signature matching [1], [18]. With the emergence of artificial intelligence and machine learning as revolutionary cybersecurity techniques, it has become possible to enhance IDS to learn from past data and adapt to new attack behaviours. Among these, ensemble learning approaches have become particularly effective because they can combine multiple poor classifiers to raise prediction accuracy. Modern organisations face greater risks of cyberattacks that range from sophisticated zero-day exploits to denial-of-service attacks due to increasing dependence on dispersed information systems, cloud infrastructures, and networked devices. Developing intelligent and automated cybersecurity systems capable of identifying threats in real time has become imperative.

Machine learning-based IDS provides higher adaptability by identifying both benign and malicious activity and learning behavioural patterns. Ensemble learning models have shown good performance among the various categories of machine learning approaches, as they may aggregate several weak learners to reduce variance, increase generalisation, and handle difficult classification tasks. The NSL-KDD dataset is a well-known benchmark for intrusion detection research. It enhances the original KDD'99 dataset by eliminating redundant and duplicate instances and makes it better suited for repeatable machine learning experiments [1].

This goal is in line with the more general objective of enhancing the precision, flexibility, and dependability of automated intrusion detection systems through the use of machine learning approaches that can successfully identify various types of assaults.

This research is guided by several key objectives toward fulfilling this stated purpose. First, multi-class classification preprocessing of the NSL-KDD dataset involves managing categorical features, scaling numerical features, and overcoming class imbalances to ensure data compatibility and enhance model performance. Following that, ensemble machine learning models, such as Random Forest and Gradient Boosting in particular, have been created and examined because of their established stability, strong generalisation capabilities, and suitability for intrusion detection in structured network environments. The systematic analysis of the performance of such models using a variety of metrics, such as accuracy, precision, recall, F1-Score, and ROC-AUC, is later conducted so as to identify relative strengths and weaknesses associated with each [2]. In addition, feature importance analysis is performed in order to determine significant network traffic features that support intrusion detection, thus enhancing interpretability and widening understanding of model decision-making

processes. Based on such findings, suggestions and insights are developed to support the development of more effective intrusion detection systems, citing the most appropriate modelling techniques, practical features, and potential ways to improve.

This paper identifies the traffic patterns in the NSL-KDD dataset and evaluates their effectiveness for intrusion detection. The research will be using two ensemble learning models, namely Random Forest and Gradient Boosting.

2. Problem Statement

Cyberattacks are increasing in scope, variety, and complexity [1]. Problems with existing IDS solutions include high false positive rates, lack of scalability, and limited flexibility. Machine learning models can be trained on benchmark datasets, such as NSL-KDD, in support of early threat identification, but effectiveness is dependent on model selection, preprocessing, and feature understanding [6]. Therefore, the aim of this work is to explore whether ensemble models can provide better performance and insights for intelligent IDS design.

3. Related Work

Intrusion detection systems have evolved from traditional rule-based and signature-driven processes to intelligent analytical systems that are capable of identifying previously unknown threats and understanding the behavioral patterns [1], [18]. Due to the fact that most of the previous systems relied on static databases of known attack signatures, they were less capable of detecting advanced persistent threats, polymorphic malware, and zero-day assaults. Due to the expanded scope and complexity, machine learning and data-driven approaches became feasible replacements to automatically detect anomalies and make inferences from the previously seen traffic.

To overcome redundancy problems in the previous KDD-99 dataset, studies like Tavallaei et al. (2009) established the NSL-KDD dataset 13. Since Random Forests are robust and able to handle high-dimensional data, studies have shown their utility for intrusion detection. Moreover, gradient boosting models have been found useful, particularly in identifying invasions in the minority class.

Ensemble learning algorithms, having resilient, stable performance with noisy and high-dimensional network traffic data, have gained increasing popularity in recent research. The bagging-based ensemble algorithm Random Forest has gained a reputation for being very robust to noise and able to model complex, non-linear incursion behavior easily. Several researchers have demonstrated that when applied to NSL-KDD, RF has consistently outperformed many traditional classifiers such as logistic regression, decision trees, and naïve Bayes. The technique is suitable for IDS scenarios due to its internal process of feature selection and handling mixed data types. Boosting-based approaches, particularly Gradient Boosting and its enhanced variant XGBoost, have also received a great deal of attention parallel to RF due to their ability to yield highly accurate models by means of iterative learning [5], [11], [14]. Various studies have indeed illustrated that in the case of problematic incursion classes such as R2L and U2R, which are often poorly represented and challenging to classify, Gradient Boosting models tend to achieve higher detection rates. Moreover, boosting algorithms include regularisation and error-driven tree construction in order to enhance learning with respect to low-frequency attack patterns.

However, the research now in publication identifies a number of gaps. First, a lot of research assesses ensemble models without using uniform preprocessing techniques, which makes comparison challenging because categorical encoding, feature scaling, and class imbalance handling vary throughout trials. Second, fewer studies offer a controlled, side-by-side evaluation employing similar feature engineering and model evaluation pipelines, even though Random Forest and Gradient Boosting have both been researched separately. Additionally, although feature importance analysis is recognised as helpful, it is frequently ignored or only briefly described, which restricts interpretability and the capacity to obtain practical insights for IDS design.

The current study uses a consistent preprocessing approach that includes stratified sampling, numerical scaling, and category encoding in order to close these gaps. This processed dataset

was used to train and fine-tune Random Forest and Gradient Boosting, which were then assessed using a variety of metrics, including accuracy, precision, recall, F1-score, confusion matrices, and computational cost. To ascertain which traffic characteristics have the most impact on model performance, the method also incorporates feature importance analysis. This study contributes to ongoing research in machine-learning-based intrusion detection by offering empirical insights into performance trade-offs, computational efficiency, and interpretability through a controlled experimental setting comparison of both ensembles.

4. Methodology

Using the NSL-KDD dataset to create, assess and contrast ensemble-based intrusion detection models using an organised machine learning methodology. The five main phases of the methodology are dataset gathering, preprocessing, model creation, model evaluation, and interpretation.

4.1 Acquisition of Dataset

Because of its extensive use in intrusion detection research and its enhancements over the original KDD'99 dataset, such as fewer duplicate entries and a more evenly distributed class, the NSL-KDD dataset was chosen [13]. The original benchmark evaluation structure was maintained by loading the dataset into two distinct splits: a training file (kdd_train.xlsx) and a testing file (kdd_test.xlsx).

4.2 Data Preprocessing

To get the dataset ready for machine learning, preprocessing is crucial. Initially, a rule-based function was used to map the original dataset labels into five attack categories: Normal, DoS, Probe, R2L, and U2R. This change is consistent with earlier studies and facilitates multi-class intrusion detection.

The numerical features were then standardised using StandardScaler to normalise feature ranges, and the categorical attributes (protocol_type, service, and flag) were encoded using One-Hot Encoding. To preserve data integrity, missing values from numeric conversion were imputed with zero.

A stratified random sampling of 20,000 training rows was carried out to provide sufficient representation across all attack categories due to computational limitations and class imbalance.

Fig. 1 – Map attack counts for train and test datasets

```
Mapped class counts (train):
attack_category
normal      67343
dos         45927
probe       11656
r2l          995
u2r           52
Name: count, dtype: int64
Mapped class counts (test):
attack_category
normal      11245
dos         8095
probe       2157
r2l          918
other        100
u2r           29
Name: count, dtype: int64
```

4.3 Model Development

For comparison, two ensemble models were chosen:

1. Random Forest Classifier (bagging technique)
2. Gradient Boosting Classifier (boosting technique)

To guarantee constant transformation throughout training and testing, the models and the preprocessing stages were integrated into a scikit-learn pipeline.

The Random Forest model was subjected to hyperparameter optimization using RandomizedSearchCV, with the F1-score as the optimization metric and a stratified K-Fold cross-validation technique ($k=3$). For comparative baseline analysis, Gradient Boosting was trained with default parameters.

4.4 Model Assessment

The reserved test set was used to evaluate the model's performance. The metrics listed below were calculated:

- Accuracy
- Precision
- Recall
- F1-Score
- Confusion Matrix

Particularly for minority attack classes like R2L and U2R, these measures offer a thorough grasp of each model's multi-class classification capability. In order to compare computational efficiency, training time was also noted.

4.5 Analysis of Feature Importance

To determine which features have the greatest influence on classification decisions, feature importance values were taken from both trained models. To comprehend attack-relevant network information (such as `src_bytes`, `dst_bytes`, `flag`, and `diff_srv_rate`), a sorted priority table and visual bar charts were created.

4.6 Interpretation and Creation of Insight

In order to assess trade-offs between classification accuracy, computational efficiency, and interpretability, the last step involved analysing model findings. Recommendations for enhancing IDS design and guiding future work were made using the insights.

5. Results and Discussion

The experimental findings from training and assessing the Random Forest and Gradient Boosting models on the NSL-KDD dataset are shown in this section. To determine the relative advantages and disadvantages of the two ensemble techniques, the results are analysed in terms of classification performance, computational efficiency, and feature importance.

5.1 Classification Performance

On the NSL-KDD dataset, both ensemble models performed well; nonetheless, there were noticeable differences between particular attack categories. The Random Forest model's overall accuracy was 0.9192, whereas the Gradient Boosting model's was 0.9308. This implies that the dataset's complex decision boundaries were better learned by Gradient Boosting.

Both models performed remarkably well on Normal and DoS traffic, with precision and recall values above 0.98 in the majority of cases, according to a thorough analysis of the classification report. This is consistent with the dataset's class distributions, which show that these groups are the most prevalent. But in the minority courses, especially R2L and U2R, where sample sizes were much smaller, performance sharply decreased.

Table I – Class-Level Model Performance

Class	RF F1-Score	GB F1-Score	Best Model
DoS	0.9617	0.9737	GB
Normal	0.9274	0.9398	GB
Probe	0.9106	0.9052	RF
R2L	0.2327	0.3612	GB
U2R	0.1765	0.2740	GB
Other	0.0000	0.0000	--

5.2 Interpretation of the Confusion Matrix

Classification behaviour is further demonstrated by the confusion matrices. Both models exhibit consistent prediction accuracy for Normal, DoS, and Probe assaults, as well as high diagonal dominance for the majority classes. R2L and U2R were often projected as Normal or Probe traffic, and misclassification patterns were more concentrated in minority groups. This demonstrates how class imbalance in intrusion detection remains a problem even after stratified sampling.

Fig. 2 – Random Forest Confusion Matrix





Fig. 3 – Gradient Boosting Confusion Matrix



5.3 Efficiency of Computation

There was a notable variation in training time when computational costs were compared. Under the identical circumstances, the Random Forest model finished training in 0.40 seconds, whereas Gradient Boosting took 15.58 seconds. Gradient Boosting required significantly more computing power, even if it delivered better minority class handling and higher accuracy [4], [5], [11]. This trade-off suggests that Gradient Boosting is a better option for situations demanding higher detection accuracy, whereas Random Forest may be more appropriate for deployment where quick retraining or resource constraints are a requirement.

5.4 Feature Importance Analysis

Consistent high-impact characteristics were found in both models through feature importance analysis. `src_bytes`, `same_srv_rate`, `diff_srv_rate`, `dst_host_serror_rate`, and `logged_in` were the most significant features. These characteristics have been connected to questionable connection patterns and unusual session behaviour in the past, demonstrating their function in separating legitimate activity from malicious traffic. The efficiency of one-hot encoded categorical variables in aiding intrusion detection is demonstrated by the presence of service- and flag-encoded features (e.g., `flag_S0`, `flag_SF`) among top contributors.

Fig. 4 – Top Feature Importance

Top features by RF importance:			
	feature	rf_importance	gb_importance
0	<code>src_bytes</code>	0.100335	0.387996
1	<code>diff_srv_rate</code>	0.066176	0.031542
2	<code>flag_S0</code>	0.059355	0.002485
3	<code>dst_host_serror_rate</code>	0.057830	0.093042
4	<code>flag_SF</code>	0.057699	0.000022
5	<code>dst_host_diff_srv_rate</code>	0.057490	0.051138
6	<code>dst_bytes</code>	0.049812	0.025881
7	<code>same_srv_rate</code>	0.047991	0.165326
8	<code>dst_host_same_srv_rate</code>	0.035741	0.012401
9	<code>serror_rate</code>	0.034991	0.022129
10	<code>count</code>	0.033569	0.020397
11	<code>dst_host_srv_diff_host_rate</code>	0.032294	0.000577
12	<code>dst_host_srv_count</code>	0.029690	0.001365
13	<code>dst_host_same_src_port_rate</code>	0.029662	0.007624
14	<code>srv_serror_rate</code>	0.028490	0.000102
15	<code>logged_in</code>	0.022835	0.001407
16	<code>srv_count</code>	0.021627	0.000044
17	<code>dst_host_srv_rerror_rate</code>	0.020713	0.000154
18	<code>dst_host_count</code>	0.020206	0.000454
19	<code>dst_host_srv_serror_rate</code>	0.019995	0.000144


5.5 Overview of the Discussion


The findings show that both ensemble models are useful for IDS tasks; however, Gradient Boosting consistently performed better than Random Forest on important performance parameters, especially for attack types that are underrepresented. Class imbalance is still a

major problem despite this increase, and future improvement techniques like oversampling, cost-sensitive learning, or synthetic data (like SMOTE) may boost performance even more.

Overall, the results verify that ensemble techniques, particularly boosting, offer a benefit in intrusion detection applications. Feature importance analysis provides comprehensible information that can help with system optimisation and direct security analysis in setting priorities for characteristics important to network threat detection.

Fig. 5 – Feature Importance Graph of Random Forest and Gradient Boosting





6. Conclusion and Recommendations

Using the NSL-KDD dataset, this study explained the efficacy of ensemble-based machine learning models for multi-class intrusion detection. In order to provide practical insights for enhancing intrusion detection systems, the study sought to preprocess the dataset, create and assess ensemble classifiers, compare their performance using common evaluation metrics, and analyse feature importance.

With both Random Forest and Gradient Boosting obtaining good classification performance, the results show that ensemble learning techniques are useful for intrusion detection [2], [8], [11]. With an overall accuracy of 93.08%, Gradient Boosting outperformed Random Forest, which came in at 91.92%. The thorough analysis revealed that while both models did remarkably well in identifying the majority attack classes like Normal, DoS, and Probe, they performed poorly for minority classes like R2L and U2R, which are still difficult to identify because of a significant class imbalance and few discernible behavioural patterns.

In terms of computational efficiency, Random Forest showed a far faster training time (0.40 seconds) than Gradient Boosting (15.58 seconds), indicating that the deployment environment's computational limitations and retraining requirements may influence the model selection. Attributes like `src_bytes`, `diff_srv_rate`, `dst_host_serror_rate`, and `flag` values contributed the most to model decision-making, according to feature importance analysis, highlighting their significance in differentiating between malicious network behaviour and genuine traffic.

In general, the results show that ensemble learning has a great deal of promise to improve automated intrusion detection, especially when boosting-based techniques are used. Nonetheless, there are still issues with class distribution management and low-frequency attack detection.

6.1 Recommendations

A number of suggestions are made in light of the study's conclusions to help enhance machine learning-based intrusion detection systems in the future. In order to increase the detection accuracy of minority attack categories like R2L and U2R, future trials should use methods like SMOTE, adaptive resampling, or cost-sensitive learning. Addressing class imbalance is still a top concern. To maximise model performance, additional hyperparameter adjustment is also required, especially for boosting-based algorithms where more thorough optimization employing contemporary variations like XGBoost, LightGBM, or CatBoost may improve both predictive power and dimensionality reduction techniques, including principal component analysis or mutual information filtering, should also be taken into account because some features were shown to have little effect and eliminating them could increase model efficiency without. For instance, to capture sequential dependencies and intricate behavioural patterns in network traffic, future research may also benefit from hybrid and deep learning-based approaches that combine ensemble techniques with neural network architectures like convolutional neural networks (CNNs) or long short-term memory (LSTM) models. Real-time deployment testing should be given priority over offline evaluation in

order to evaluate the scalability, latency, and reliability of the model in real-world or simulated network scenarios. Lastly, in operational cybersecurity situations, using explainable AI frameworks like SHAP or LIME will improve model transparency, facilitate analyst interpretation, and boost confidence in automated decision-making.

This investigation adds to the increasing amount of research that shows how effective ensemble machine learning techniques are for intrusion detection. The paper offers useful insight into the trade-offs between accuracy, computational cost, and interpretability by comparing Random Forest with Gradient Boosting under a single preprocessing and assessment framework. Ensemble models have great potential to enable intelligent, adaptive, and automated cybersecurity defence systems with further optimisation and class-balancing techniques.

6.2 Limitations

Even though both ensemble models show encouraging performance, it is important to recognise several of the study's shortcomings. First, although being extensively utilised in scholarly research, the NSL-KDD dataset is thought to be out of date and does not accurately capture the complexity of contemporary network settings, including encrypted traffic, IoT communication patterns, cloud infrastructures, and changing threat behaviours [1], [13]. Consequently, the findings have limited applicability to modern real-world systems. Second, there is a significant class imbalance in the dataset, especially for uncommon attack categories like R2L and U2R. This imbalance greatly impacted model performance and led to the misdiagnosis of minority attack types, even with stratified sampling and tuning efforts. Third, latency, throughput limitations, and real-time packet processing were not assessed because the studies were carried out offline in a controlled setting. This makes it impossible to draw firm conclusions regarding operational performance in real-world deployment settings. Furthermore, even though hyperparameter tuning was used, optimisation limitations, especially for boosting-based models, which might have stopped additional performance improvements. Lastly, the study did not investigate sophisticated feature engineering or deep temporal analysis methods that could catch intricate attack patterns; instead, it concentrated on traditional feature encoding. These drawbacks emphasise the necessity of additional testing with more varied datasets, sophisticated balancing techniques, and real-time evaluation frameworks.

7. References

- [1] K. A, G. I, V. P and K. J, “Survey of intrusion detection systems: techniques, datasets and challenges,” *Cybersecurity*, vol. 2, no. 1, pp. 1-22 , 2019.
- [2] C. Reddy and K. K. et al, “Twined ensemble framework for network security: integrating Random Forest, AdaBoost, and Gradient Boosting for enhanced intrusion detection.,” *Discover Internet of Things*, vol. 1, no. 5, p. 107, 2025.
- [3] V. A and R. V, “Machine learning based intrusion detection systems for IoT applications,” *Machine learning based intrusion detection systems for IoT applications*, vol. 111, p. 2287–2310, 2020.
- [4] Z. S, B. A. N, K. S and S. O. T, “Random forests approach for network intrusion detection system,” *Journal of Information Security and Applications*, vol. 54, pp. 1-10, 2020.
- [5] A. E. A. M, “A deep learning-based cybersecurity intrusion detection system using XGBoost,” *Applied Sciences*, vol. 10, no. 14, pp. 1-15, 2020.
- [6] F. G, F. J, d. A. L and M. J, “Feature selection techniques for intrusion detection systems: a review,” *Journal of Information and Communication Technology*, vol. 18, no. 3, pp. 419-440, 2019.
- [7] S. C and Y. W, “A Machine Learning Dataset for Real-Time Intrusion Detection on Industrial Control Systems,” *International Journal of Critical Infrastructure Protection*, vol. 36, no. 100464, 2022.
- [8] V. A, B. D and G. S. S, “Anomaly Detection in Large-Scale Cybersecurity Graphs using Ensemble Learning,” *Computers and Security*, vol. 98, no. 102017, 2020.
- [9] T. R. M, S. S and G. C, “Why Should I Trust You?': Explaining the Predictions of Any Classifier,” *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135-1144, 2016.
- [10] L. J and I. L. S, “A Unified Approach to Interpreting Model Predictions,” *Proc. 31st Conference on Neural Information Processing Systems (NeurIPS)*, pp. 4765-4774, 2017.
- [11] T. A, M. N and M. K. N, “Hybrid Deep Learning and Ensemble Method for Intrusion Detection,” *Journal of Information Security and Applications*, vol. 57, no. 102710, 2021.

- [12] B. B and R. F, "Wild Patterns: Ten Years After the Rise of Adversarial Machine Learning," *Pattern Recognition*, vol. 84, pp. 317-331, 2018.
- [13] C. G and H. J, "Generation of a New IDS Test Dataset: Time-based Contiguous Network Traffic," *Proc. 4th International Conference on Security of Information and Networks (SIN)*, pp. 192-199, 2011.
- [14] W. H , W. J and Z. Z, "Network Intrusion Detection Using LightGBM with Feature Engineering," *Security and Communication Networks*, vol. 2019, no. 8712845, 2019.
- [15] K. K, S. K. V, B. K. P. L and C. J. R, "An Improved Random Forest for Intrusion Detection using Gini Index," *International Journal of Computer Applications*, vol. 128, no. 4, pp. 28-34, 2015.
- [16] N. T. T, X. B, Y. L and Y. R, "Real-time Anomaly-based Network Intrusion Detection Platform with Automated Feature Extraction and Learning," *Computers and Security*, vol. 78, pp. 316-332, 2018.
- [17] F. M, F. K. M, H. A and R. A, "Performance Comparison of Machine Learning Techniques for Intrusion Detection in Networking Environment," *Journal of Information and Communication Technology*, vol. 14, no. 2, pp. 149-164, 2015.
- [18] C. V, B. A and K. V, "Anomaly Detection: A Survey," *ACM Computing Surveys*, vol. 41, no. 3, 2009.
- [19] C. V, W. B. K, O. H. L and P. K. W, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321-357, 2002.