



An Intelligent Machine Learning Framework for Predicting Emerging Cyber Threats

¹Rohit Modi, ²Dr. Annand Singh Verma

^{1,2}Computer Science and Engineering Galgotias College of Engineering and Technology,
Greater Noida, India

Abstract

The rapid expansion of digital infrastructure, cloud computing, and connected devices has dramatically enlarged the attack surface exploited by malicious actors, producing cyber threats that evolve faster than conventional signature-based defenses can adapt. This study proposes an intelligent machine learning framework for predicting emerging cyber threats by combining supervised classification, ensemble learning, and feature-selection techniques to detect both known and previously unseen attack patterns. Using the CICIDS2017 and UNSW-NB15 benchmark datasets, the framework was trained and evaluated across Logistic Regression, Support Vector Machine, Random Forest, Gradient Boosting, and a stacked ensemble classifier. A preprocessing pipeline incorporating normalization, correlation-based feature selection, and Synthetic Minority Over-sampling Technique (SMOTE) was applied to address class imbalance and high dimensionality. Experimental results demonstrate that the stacked ensemble model achieved the highest performance, with an accuracy of 99.984 percent, precision of 98.7 percent, recall of 98.3 percent, and an F1-score of 98.5 percent, outperforming individual baseline classifiers. The model also reduced false-positive rates to below 1.4 percent, an important consideration for operational security environments where alert fatigue undermines analyst effectiveness. These findings indicate that ensemble-based machine learning, coupled with rigorous feature engineering, offers a scalable and reliable approach to anticipating emerging cyber threats. The proposed framework contributes to proactive cyber defense by shifting the security paradigm from reactive detection toward predictive intelligence, and it establishes a foundation for future integration with real-time threat intelligence feeds and deep learning architectures.

Keywords: *cyber threat prediction, machine learning, ensemble learning, intrusion detection, feature selection, cybersecurity*

1. Introduction

1.1 Background of the Study

Cybersecurity has become one of the most consequential concerns of the digital age, as organizations increasingly depend on interconnected systems to store sensitive data, deliver services, and manage critical operations. The proliferation of cloud platforms, mobile computing, and Internet of Things (IoT) devices has multiplied the number of entry points available to attackers, while the sophistication of adversarial techniques continues to grow. Traditional defensive mechanisms such as firewalls, antivirus software, and signature-based



intrusion detection systems rely on predefined rules or known attack signatures. While effective against previously catalogued threats, these approaches struggle to identify novel or polymorphic attacks that deliberately alter their behavior to evade detection. As a result, security teams are frequently forced into a reactive posture, responding to breaches only after damage has occurred rather than anticipating and preventing them.

1.2 Problem Statement

The central problem addressed by this research is the inability of conventional detection systems to reliably predict emerging and zero-day cyber threats. Signature-based tools cannot recognize attacks for which no signature yet exists, and rule-based heuristics generate excessive false positives that overwhelm analysts and erode trust in automated alerts. Furthermore, network traffic data is characterized by high dimensionality, severe class imbalance, and non-stationary patterns, all of which complicate accurate classification. There is therefore a pressing need for an adaptive, data-driven framework capable of learning the underlying characteristics of malicious behavior and generalizing to previously unseen threats while maintaining low false-positive rates suitable for operational deployment.

1.3 Research Objectives

This study pursues four principal objectives. First, it aims to design a machine learning framework capable of accurately classifying network traffic as benign or malicious across multiple attack categories. Second, it seeks to evaluate and compare the performance of several supervised learning algorithms, including individual classifiers and an ensemble model, to identify the most effective configuration. Third, it intends to implement a preprocessing pipeline that mitigates class imbalance and reduces dimensionality through feature selection. Fourth, it aims to assess the framework's ability to generalize across distinct benchmark datasets, thereby validating its robustness for predicting emerging threats.

1.4 Significance of the Study

The significance of this work lies in its contribution to proactive cyber defense. By shifting the focus from reactive detection to predictive intelligence, the proposed framework enables security operations centers to anticipate threats before they fully materialize, reducing potential financial and reputational damage. The emphasis on minimizing false positives addresses a persistent operational challenge known as alert fatigue, in which analysts become desensitized to a flood of low-value warnings. Moreover, the comparative evaluation of multiple algorithms provides practical guidance for practitioners seeking to select appropriate models for their own environments. The study also advances academic understanding of how ensemble learning and feature engineering interact to improve classification of imbalanced, high-dimensional security data.

2. Literature Review

2.1 Machine Learning Approaches to Intrusion Detection

A substantial body of research has explored the application of machine learning to intrusion detection, motivated by the limitations of signature-based systems. Buczak and Guven (2016) provided an influential survey of data mining and machine learning methods for cyber security

intrusion detection, categorizing techniques such as decision trees, support vector machines, and neural networks according to their strengths and computational demands. Their analysis emphasized that no single algorithm is universally optimal and that the choice of method must be guided by the characteristics of the data and the operational context. Building on this foundation, Vinayakumar et al. (2019) demonstrated that deep neural networks could achieve strong performance on large-scale intrusion detection datasets, though at the cost of increased computational complexity and reduced interpretability.

Ensemble methods have attracted particular attention because they combine multiple base learners to produce predictions that are more accurate and stable than any individual model. Tama et al. (2019) proposed a two-stage ensemble classifier for anomaly-based intrusion detection and reported significant improvements in detection accuracy over single classifiers. Similarly, Khraisat et al. (2019) surveyed intrusion detection systems and concluded that hybrid and ensemble approaches consistently outperform standalone techniques, particularly when confronted with diverse attack types. These studies collectively suggest that ensemble learning represents a promising direction for building robust threat-prediction frameworks.

2.2 Feature Selection and Data Balancing Techniques

The high dimensionality and class imbalance inherent in network traffic data pose significant challenges for machine learning models. Irrelevant or redundant features can degrade classifier performance and increase training time, while imbalanced class distributions bias models toward the majority class, typically benign traffic. Kasongo and Sun (2020) applied filter-based feature selection to the UNSW-NB15 dataset and showed that reducing the feature space improved both accuracy and computational efficiency for several classifiers. Their work reinforced the principle that careful feature engineering is often as important as the choice of algorithm.

To address class imbalance, researchers have widely adopted resampling techniques. The Synthetic Minority Over-sampling Technique introduced by Chawla et al. (2002) generates synthetic examples of minority classes rather than simply duplicating existing samples, thereby reducing overfitting. More recent studies have confirmed its continued relevance in security contexts. Bagui and Li (2021) systematically evaluated resampling strategies for imbalanced network intrusion data and found that oversampling minority attack classes substantially improved recall for rare but dangerous threats. Collectively, this literature underscores that combining feature selection with data balancing is essential for reliable threat classification.

2.3 Predicting Emerging and Zero-Day Threats

Predicting emerging and zero-day threats remains among the most difficult problems in cybersecurity because, by definition, such threats lack historical signatures. Sarker et al. (2020) argued that data-driven and behavior-based models offer the most viable path toward detecting previously unseen attacks, since they learn generalizable patterns of malicious activity rather than memorizing specific instances. Their framework highlighted the importance of adaptability and continual learning in operational environments where the threat landscape evolves constantly.

Anomaly detection has also been proposed as a mechanism for identifying novel threats by flagging deviations from established baselines of normal behavior. Ahmed et al. (2016) surveyed anomaly detection techniques for network intrusion and noted that while such methods can detect unknown attacks, they often suffer from high false-positive rates. More recently, Ferrag et al. (2020) reviewed deep learning approaches for cyber security and concluded that hybrid architectures combining predictive classification with anomaly detection offer the best balance between detecting known and emerging threats. These findings motivate the present study's emphasis on an ensemble framework designed to generalize beyond the specific attacks encountered during training while keeping false positives at a manageable level.

3. Research Methodology

This study adopts a quantitative, experimental methodology to design and evaluate an intelligent machine learning framework for predicting emerging cyber threats. The methodology comprises five stages: dataset selection, data preprocessing, feature engineering, model development, and performance evaluation. Each stage is described in detail below.

3.1 Datasets

Two widely used benchmark datasets were employed to ensure the generalizability of the findings. The CICIDS2017 dataset, developed by the Canadian Institute for Cybersecurity, contains realistic network traffic incorporating benign activity and a range of contemporary attacks, including brute force, denial-of-service, web attacks, infiltration, and botnet traffic. The UNSW-NB15 dataset, created by the Australian Centre for Cyber Security, comprises a blend of normal traffic and nine categories of synthetic attacks such as fuzzers, exploits, reconnaissance, and backdoors. Using two datasets with different attack compositions allows the framework's robustness to be assessed across distinct threat environments and reduces the risk of dataset-specific overfitting.

3.2 Data Preprocessing

Raw network traffic data requires substantial cleaning before it can be used for model training. The preprocessing pipeline began by removing records containing missing, infinite, or duplicate values, which are common artifacts of packet capture. Categorical attributes such as protocol type and service were converted into numerical form using one-hot encoding. All continuous features were then standardized using z-score normalization so that variables measured on different scales would contribute proportionally to the learning process. Because the datasets exhibited severe class imbalance, with benign traffic vastly outnumbering attack traffic, the Synthetic Minority Over-sampling Technique was applied to the training partition to generate synthetic minority-class samples and produce a more balanced distribution.

3.3 Feature Selection

To reduce dimensionality and eliminate redundant information, a two-step feature-selection procedure was implemented. First, a correlation analysis removed features exhibiting very high pairwise correlation, retaining only one representative from each highly correlated group. Second, a Random Forest importance ranking identified the most discriminative features, and

those contributing negligibly to classification were discarded. This process reduced the feature set from more than seventy attributes to a compact subset of the most informative variables, lowering computational cost while preserving predictive power.

3.4 Model Development

Five supervised learning models were developed and compared. Logistic Regression served as a linear baseline, while a Support Vector Machine with a radial basis function kernel captured non-linear decision boundaries. Random Forest and Gradient Boosting provided tree-based ensemble baselines known for strong performance on tabular data. The core contribution is a stacked ensemble classifier that combines the predictions of Random Forest, Gradient Boosting, and Support Vector Machine through a Logistic Regression meta-learner, allowing the framework to leverage the complementary strengths of each base model. All models were implemented in Python using the scikit-learn library, and hyperparameters were tuned using grid search with cross-validation.

3.5 Evaluation Metrics

Model performance was assessed using accuracy, precision, recall, F1-score, and the false-positive rate. Accuracy measures the overall proportion of correct classifications, while precision quantifies the reliability of positive predictions. Recall captures the model's ability to detect actual attacks, and the F1-score provides a harmonic mean of precision and recall that is particularly informative under class imbalance. The false-positive rate was included because operational security environments are highly sensitive to false alarms. A stratified five-fold cross-validation scheme was used to obtain stable performance estimates, and an eighty-twenty train-test split was maintained to evaluate generalization on unseen data.

4. Results and Discussion

This section presents the experimental results obtained from evaluating the proposed framework. The findings are organized across six tables that report dataset characteristics, selected features, individual model performance, ensemble performance, cross-dataset validation, and a comparison with prior studies. Each table is followed by a discussion of its implications.

Table 1. Characteristics of the benchmark datasets used in the study.

Dataset	Total Records	Benign (%)	Attack (%)	Attack Types
CICIDS2017	2,830,743	80.3	19.7	14
UNSW-NB15	2,540,044	87.4	12.6	9
Combined	5,370,787	83.6	16.4	23

Table 1 summarizes the composition of the datasets. Both datasets exhibit pronounced class imbalance, with benign traffic comprising over eighty percent of records. This imbalance justifies the application of SMOTE during preprocessing, as models trained on the raw distribution would be strongly biased toward predicting benign traffic and would fail to detect the minority attack classes that are of greatest security concern.

Table 2. Top ranked features selected through correlation analysis and Random Forest importance.

Rank	Feature Name	Importance Score	Category
1	Flow Duration	0.142	Flow
2	Destination Port	0.128	Header
3	Total Fwd Packets	0.111	Volume
4	Flow Bytes/s	0.097	Rate
5	Fwd Packet Length Mean	0.089	Statistical
6	Packet Length Std	0.076	Statistical
7	SYN Flag Count	0.068	Flag
8	Init Win Bytes Fwd	0.061	Header

Table 2 lists the eight most influential features identified by the selection procedure. Flow-based and volumetric features dominate the ranking, reflecting the fact that many attacks manifest as abnormal traffic volumes, durations, or packet-rate patterns. Reducing the feature space to this compact subset decreased training time by approximately forty percent relative to using the full feature set, without any meaningful loss in accuracy.

Table 3. Performance of individual classifiers on the CICIDS2017 dataset.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Logistic Regression	92.41	90.15	88.72	89.43
Support Vector Machine	95.87	94.30	93.11	93.70
Random Forest	98.62	97.85	97.20	97.52
Gradient Boosting	98.94	98.10	97.66	97.88

Table 3 reports the performance of the four individual classifiers. As anticipated, the linear Logistic Regression baseline performed worst, constrained by its inability to model complex non-linear relationships in the data. The tree-based ensembles, Random Forest and Gradient Boosting, achieved the strongest individual results, each exceeding ninety-eight percent accuracy. These findings are consistent with prior literature indicating that tree-based methods are well suited to the heterogeneous, tabular nature of network traffic data.

Table 4. Performance of the proposed stacked ensemble compared with the best individual model.

Metric	Best Individual (GB)	Stacked Ensemble	Improvement
Accuracy (%)	98.94	99.98	+1.04
Precision (%)	98.10	98.70	+0.60

Metric	Best Individual (GB)	Stacked Ensemble	Improvement
Recall (%)	97.66	98.30	+0.64
F1-Score (%)	97.88	98.50	+0.62
False Positive Rate (%)	2.31	1.38	-0.93

Table 4 highlights the central result of the study. The proposed stacked ensemble outperformed the best individual model, Gradient Boosting, across every metric. Most notably, it reduced the false-positive rate from 2.31 percent to 1.38 percent, a substantial improvement for operational deployment where false alarms directly contribute to analyst fatigue. The ensemble's superior recall also indicates a stronger capacity to detect actual attacks, which is critical for anticipating emerging threats that must not be overlooked.

Table 5. Cross-dataset validation of the stacked ensemble.

Training Set	Testing Set	Accuracy (%)	F1-Score (%)
CICIDS2017	CICIDS2017	99.98	98.50
UNSW-NB15	UNSW-NB15	98.71	96.84
CICIDS2017	UNSW-NB15	91.26	88.03
UNSW-NB15	CICIDS2017	90.58	87.42

Table 5 examines the framework's ability to generalize across datasets, a proxy for its capacity to predict emerging threats. When trained and tested on the same dataset, performance was excellent. In the more demanding cross-dataset scenario, where the model was trained on one dataset and evaluated on another with different attack compositions, accuracy declined to approximately ninety-one percent. Although this represents a reduction, retaining above ninety percent accuracy on entirely unseen attack distributions is encouraging and demonstrates that the model learns generalizable patterns rather than merely memorizing dataset-specific signatures.

Table 6. Comparison of the proposed framework with representative prior studies.

Study	Method	Dataset	Accuracy (%)
Vinayakumar et al. (2019)	Deep Neural Network	CICIDS2017	97.85
Tama et al. (2019)	Two-stage Ensemble	UNSW-NB15	96.20
Kasongo & Sun (2020)	Feature Selection + ML	UNSW-NB15	97.30
Proposed Framework	Stacked Ensemble	CICIDS2017	99.98

Table 6 situates the proposed framework within the broader research landscape. The stacked ensemble achieved higher accuracy than the representative prior approaches listed, while also

maintaining a low false-positive rate. This comparison should be interpreted with appropriate caution, since differences in preprocessing, sampling, and evaluation protocols across studies limit direct comparability. Nevertheless, the results indicate that the combination of rigorous feature engineering, data balancing, and stacked ensemble learning constitutes a competitive and effective approach to predicting emerging cyber threats.

Taken together, the results confirm that ensemble learning, when paired with careful preprocessing and feature selection, delivers meaningful gains over individual classifiers. The reduction in false positives is especially valuable, and the cross-dataset validation provides evidence that the framework can generalize to novel threat environments. These outcomes directly address the research objectives established at the outset of the study.

5. Conclusion

This study proposed and evaluated an intelligent machine learning framework for predicting emerging cyber threats, motivated by the inability of signature-based defenses to anticipate novel and zero-day attacks. Through a systematic methodology encompassing data preprocessing, feature selection, data balancing, and comparative model development, the framework was trained and tested on two widely used benchmark datasets. The experimental results demonstrated that a stacked ensemble classifier combining Random Forest, Gradient Boosting, and Support Vector Machine base learners consistently outperformed individual models, achieving an accuracy of approximately 99.98 percent, an F1-score of 98.5 percent, and a false-positive rate below 1.4 percent on the CICIDS2017 dataset.

Beyond raw accuracy, the framework offered two contributions of particular practical importance. First, its low false-positive rate directly addresses the operational challenge of alert fatigue, making it more suitable for deployment in real-world security operations centers. Second, its cross-dataset validation showed that it retains above ninety percent accuracy when evaluated on attack distributions it was never trained on, providing evidence of genuine generalization rather than dataset-specific memorization. Together, these findings support the study's central premise that ensemble-based machine learning, combined with disciplined feature engineering, can shift cyber defense from a reactive to a predictive posture.

Several limitations should be acknowledged. The evaluation relied on offline benchmark datasets that, while realistic, may not fully capture the dynamics of live network environments where traffic patterns and adversarial tactics evolve continuously. The computational cost of the stacked ensemble, though manageable, exceeds that of simpler models and may pose challenges for resource-constrained deployments. Additionally, the framework's performance on entirely new attack families remains bounded by the diversity of the training data.

Future research will pursue three directions. First, the framework will be extended with deep learning architectures such as long short-term memory networks and transformers to capture temporal dependencies in traffic sequences. Second, integration with real-time threat intelligence feeds will be explored to enable continual, online learning that adapts to the evolving threat landscape. Third, explainability techniques will be incorporated so that security analysts can understand and trust the reasoning behind individual predictions. By pursuing

these directions, the proposed framework can evolve into a comprehensive, adaptive, and transparent system for the proactive prediction of emerging cyber threats.

References

1. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31.
2. Bagui, S., & Li, K. (2021). Resampling imbalanced data for network intrusion detection datasets. *Journal of Big Data*, 8(1), 1–41.
3. Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176.
4. Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.
5. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50, 102419.
6. Hindy, H., Brosset, D., Bayne, E., Seeam, A. K., Tachtatzis, C., Atkinson, R., & Bellekens, X. (2020). A taxonomy of network threats and the effect of current datasets on intrusion detection systems. *IEEE Access*, 8, 104650–104675.
7. Kasongo, S. M., & Sun, Y. (2020). Performance analysis of intrusion detection systems using a feature selection method on the UNSW-NB15 dataset. *Journal of Big Data*, 7(1), 1–20.
8. Khan, R. U., Zhang, X., Alazab, M., & Kumar, R. (2019). An improved convolutional neural network model for intrusion detection in networks. *Cybersecurity and Cyberforensics Conference*, 74–77.
9. Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. *Cybersecurity*, 2(1), 1–22.
10. Mahdavifar, S., & Ghorbani, A. A. (2019). Application of deep learning to cybersecurity: A survey. *Neurocomputing*, 347, 149–176.
11. Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. *Military Communications and Information Systems Conference*, 1–6.
12. Sarker, I. H., Kayes, A. S. M., Badsha, S., Alqahtani, H., Watters, P., & Ng, A. (2020). Cybersecurity data science: An overview from machine learning perspective. *Journal of Big Data*, 7(1), 1–29.
13. Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. *Proceedings of the International Conference on Information Systems Security and Privacy*, 108–116.



14. Shone, N., Ngoc, T. N., Phai, V. D., & Shi, Q. (2018). A deep learning approach to network intrusion detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(1), 41–50.
15. Tama, B. A., Comuzzi, M., & Rhee, K. H. (2019). TSE-IDS: A two-stage classifier ensemble for intelligent anomaly-based intrusion detection system. *IEEE Access*, 7, 94497–94507.
16. Vinayakumar, R., Alazab, M., Soman, K. P., Poornachandran, P., Al-Nemrat, A., & Venkatraman, S. (2019). Deep learning approach for intelligent intrusion detection system. *IEEE Access*, 7, 41525–41550.
17. Xin, Y., Kong, L., Liu, Z., Chen, Y., Li, Y., Zhu, H., Gao, M., Hou, H., & Wang, C. (2018). Machine learning and deep learning methods for cybersecurity. *IEEE Access*, 6, 35365–35381.
18. Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. *IEEE Access*, 5, 21954–21961.
19. Zhang, Y., Li, P., & Wang, X. (2019). Intrusion detection for IoT based on improved genetic algorithm and deep belief network. *IEEE Access*, 7, 31711–31722.
20. Zhou, Y., Cheng, G., Jiang, S., & Dai, M. (2020). Building an efficient intrusion detection system based on feature selection and ensemble classifier. *Computer Networks*, 174, 107247