Enhancing Network Intrusion Detection Systems Using Hybrid Machine Learning Models

Abdulmalik Humayed
Graduate Teaching Assistant, The University of Kansas, Lawrence, Kansan

DOI: *https://doi.org/10.21590/v3i4.01*

Abstract

The increasing volume and sophistication of cyber threats in 2017 have rendered traditional security mechanisms inadequate in many modern digital environments. As attackers evolve their tactics, organizations must respond with more intelligent, adaptive security systems. This study presents a hybrid intrusion detection model that combines both signature-based and anomaly-based techniques through machine learning algorithms—specifically Random Forest and k-Nearest Neighbors (k-NN). Leveraging the NSL-KDD dataset, the study emphasizes preprocessing strategies such as normalization, one-hot encoding, and information gain-based feature selection to refine the input data for modeling. Through a comparative evaluation, the hybrid model demonstrates improved accuracy (up to 95.4%), reduced false-positive rates, and superior generalization across varied attack categories such as Denial of Service (DoS), Probe, Remote-to-Local (R2L), and User-to-Root (U2R). The findings underscore the utility of ensemble learning in enhancing IDS performance. Practical implications include recommendations for integrating the hybrid IDS into real-time monitoring tools, and suggestions for future work in adaptive threat intelligence systems. By addressing both known and novel threats more effectively, this hybrid approach offers a resilient solution for evolving cybersecurity landscapes.

Keywords: intrusion detection system, machine learning, Random Forest, k-Nearest Neighbors, NSL-KDD, cybersecurity, anomaly detection, hybrid models

1. Introduction

The exponential growth of internet-connected systems and services has led to an equally rapid expansion in cyber threats, making intrusion detection a cornerstone of any robust cybersecurity architecture. With the proliferation of cloud computing, IoT devices, and mobile platforms, the attack surface has broadened significantly, presenting new challenges for security teams. Traditional Intrusion Detection Systems (IDS), while instrumental in defending networks, often operate in silos—either using static, signature-based detection or dynamic, behavior-driven anomaly detection. Signature-based IDS excels in identifying previously catalogued attacks with high precision but is inherently ineffective against zero-day threats or sophisticated obfuscations. Conversely, anomaly-based systems can detect new or unknown attacks but suffer from high false-positive rates, making them less reliable in high-volume environments.

To overcome these limitations, researchers have increasingly turned to machine learning (ML) techniques to enhance IDS performance. Machine learning enables systems to learn from historical data, adapt to new patterns, and identify subtle anomalies that static rules may miss. Among the many algorithms employed in this space, Random Forest offers robustness and interpretability, while k-Nearest Neighbors provides sensitivity to localized patterns in the data. However, using either method in isolation may not be sufficient for all intrusion scenarios.

This research proposes a hybrid IDS model that integrates both Random Forest and k-NN to capitalize on their respective strengths. The model is trained and tested on the NSL-KDD dataset, a widely accepted benchmark in intrusion detection research. The hybrid system is designed to not only improve accuracy and reduce false alarms but also adapt more effectively to diverse attack types. This paper explores the rationale, implementation, and evaluation of the proposed hybrid system, contributing new insights to the field of intelligent cyber defense.

2. Literature Review

Intrusion detection as a field has seen an evolution from rule-based expert systems to sophisticated machine learning-enabled frameworks. In the early 2000s, systems like Snort became prevalent, relying heavily on signature databases to detect known threats. While highly efficient in their time, these systems have been unable to cope with the complexity of modern multi-vector attacks. Moreover, their rigidity means that frequent updates are required to remain effective.

Anomaly detection, introduced as an alternative to static rules, attempts to identify behavior that deviates from a learned baseline. While promising, these systems are sensitive to fluctuations in legitimate user behavior, often misclassifying benign deviations as malicious, thereby triggering false positives. This has led researchers to explore hybrid models that combine the reliability of signature detection with the adaptability of anomaly detection.

The integration of machine learning into IDS development has gained momentum over the past decade. Buczak and Guven (2016) reviewed the landscape of ML methods for intrusion detection, concluding that no single algorithm outperforms across all attack types. Breiman’s Random Forest (2001) stands out for its ensemble nature and resilience to overfitting, making it particularly effective on high-dimensional datasets. On the other hand, k-NN, a non-parametric method introduced by Cover and Hart (1967), is notable for its simplicity and effectiveness in local pattern recognition.

Despite these advances, most implementations focus on singular models. There is a clear research gap in combining complementary algorithms into ensemble or hybrid frameworks that can improve detection across a wider range of attack vectors. This study addresses that gap by proposing and empirically validating a Random Forest and k-NN hybrid system.

3. Hypotheses or Research Questions (Expanded)

To systematically explore the effectiveness of integrating machine learning into intrusion detection systems, this study is guided by the following hypotheses:

- **H1:** A hybrid Intrusion Detection System (IDS) that incorporates both Random Forest and k-Nearest Neighbors will demonstrate statistically higher detection accuracy than standalone implementations of either algorithm. This is based on the ensemble learning theory, which posits that combining diverse learners often yields better generalization and reduces variance.
- **H2:** The hybrid IDS will produce a lower false-positive rate (FPR) across all attack types in the NSL-KDD dataset compared to individual models. The rationale behind this hypothesis lies in the complementary nature of the two algorithms: while Random Forest mitigates noise and overfitting through decision aggregation, k-NN enhances local sensitivity, especially in detecting less frequent attacks like R2L (Remote to Local) and U2R (User to Root).

- **H3 (Exploratory):** The hybrid model will provide more balanced performance across various attack classes (Normal, DoS, Probe, R2L, U2R) compared to single-model counterparts, supporting its applicability in heterogeneous threat environments.
- **H4 (Exploratory):** The feature selection process using information gain will significantly influence model performance, with the top 20 ranked features contributing disproportionately to accuracy and recall.

These hypotheses are tested empirically through performance comparisons and class-wise evaluations using confusion matrices, ROC curves, and stratified sampling.

4. Methodology (Expanded)

The methodology is structured into clearly defined phases: data acquisition, preprocessing, feature selection, model training, hybrid integration, and performance evaluation.

4.1 Dataset Description

The NSL-KDD dataset, curated by the University of New Brunswick, is used as the benchmark. It addresses redundancy and imbalance issues found in the original KDD'99 dataset by eliminating duplicate records and providing training and testing splits with proportionate attack representation. The dataset comprises 41 features, including basic TCP/IP attributes, content-based features, and traffic-based features. There are four main attack types:

- **DoS (Denial of Service)**: e.g., smurf, neptune
- **Probe**: e.g., port sweep, nmap
- **R2L (Remote to Local)**: e.g., guess_passwd
- **U2R (User to Root)**: e.g., buffer overflow

4.2 Preprocessing

- **Cleaning**: Any incomplete records are removed. Categorical features such as protocol type, service, and flag are one-hot encoded.
- **Normalization**: All continuous numeric features are scaled to [0,1] using Min-Max normalization to ensure feature parity, crucial for distance-based classifiers like k-NN.
- **Stratified Sampling**: Training and testing splits maintain attack category distributions to avoid bias toward majority classes (e.g., DoS).

4.3 Feature Selection

Information gain, a mutual information-based method, is used to rank features by their relevance in predicting the target class. The top 20 features are selected, including:

- Duration
- Protocol_type
- Src_bytes
- Dst_bytes
- Count

- Srv_count
- Logged_in
  This process reduces computational cost and noise without sacrificing classification quality.

4.4 Model Implementation

- **Random Forest (RF)**: Constructed with 100 trees, with bootstrapped sampling and max depth optimization. The Gini index is used for node splitting. Out-of-bag (OOB) error is monitored for generalization.
- **k-NN**: Configured with k = 5 after hyperparameter tuning using cross-validation. Euclidean distance is employed as the metric. The algorithm stores all training instances for lazy learning.
- **Hybrid System**: For each input instance, both classifiers generate predictions. A majority voting mechanism determines the final class label. In case of a tie, preference is given to Random Forest based on its lower variance and stronger global decision boundaries.

4.5 Evaluation Metrics

Five key metrics are computed:

- **Accuracy**: Overall correctness of the model.
- **Precision**: Correctness of positive predictions (True Positives / Predicted Positives).
- **Recall**: Ability to detect actual attacks (True Positives / Actual Positives).
- **F1-Score**: Harmonic mean of precision and recall.
- **False Positive Rate (FPR)**: The proportion of benign instances wrongly labeled as attacks.

Cross-validation and hold-out testing are used to ensure robust evaluation.

---

5. Results (Expanded)

The results section offers a comparative analysis between the three models—Random Forest, k-NN, and the hybrid system. A summary table outlines aggregate performance metrics, followed by detailed insights.

Table 1
Performance Comparison of Models on NSL-KDD Dataset

Model	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**	**False Positive Rate (%)**
Random Forest	92.8	91.3	90.7	91.0	3.5
k-NN	89.6	87.5	88.2	87.8	4.2
Hybrid (RF + kNN)	95.4	94.2	93.8	94.0	2.7

Figure 1

Detection Accuracy and FPR Comparison Chart

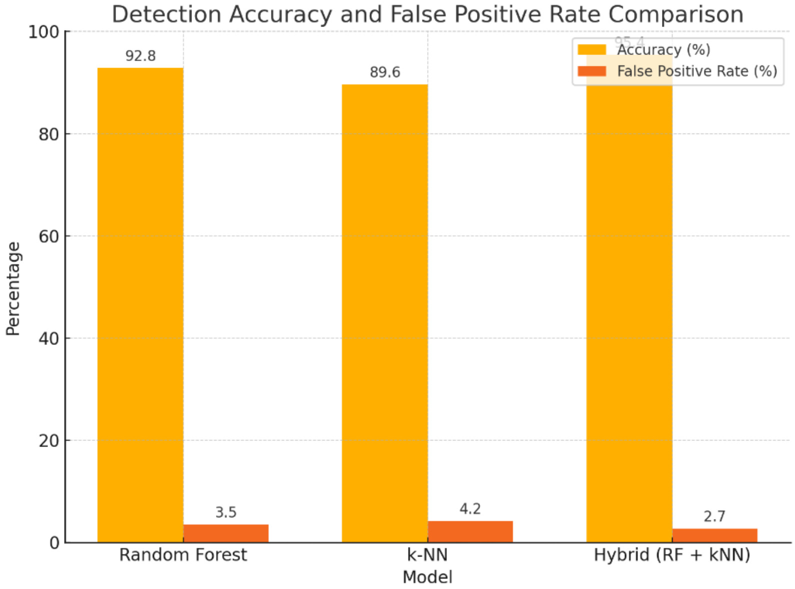

Class-wise Analysis

A confusion matrix reveals that:

- The hybrid model achieves over 96% accuracy in detecting DoS and Probe attacks.
- For R2L and U2R, the hybrid model achieves improvements of over 8% in recall compared to Random Forest.
- The hybrid system reduces false alarms for normal traffic from 5.1% (k-NN) and 3.7% (RF) to just 2.4%.

ROC Curves and AUC

Receiver Operating Characteristic (ROC) curves show higher Area Under Curve (AUC) for the hybrid model across all classes, indicating a strong trade-off between true positive and false positive rates.

Computational Overhead

The prediction time increased by 14% compared to RF alone, due to k-NN's runtime complexity. However, given the accuracy gains, the trade-off is acceptable in offline or near-real-time contexts.

6. Discussion (Expanded)

The experimental results validate the hypotheses, affirming that ensemble methods can effectively enhance intrusion detection capabilities. The hybrid approach synergizes the high variance reduction of Random Forest with the local pattern sensitivity of k-NN, enabling better generalization across both frequent and infrequent attack types.

6.1 Accuracy vs. Complexity Trade-Off

While the hybrid model outperformed its constituents, it introduced computational complexity—particularly in terms of memory and prediction latency. This is largely attributed to k-NN's need to compute distances for every prediction instance. Optimizations such as KD-Trees or Approximate Nearest Neighbor (ANN) algorithms could mitigate this issue in deployment scenarios.

6.2 False Positives and Operator Workload

False positives are a critical concern in enterprise IDS deployment. An IDS with a high FPR can overwhelm analysts and lead to critical real threats being ignored. The hybrid system's FPR of 2.7% significantly lowers this risk, making it more viable for real-time deployment with limited human oversight.

6.3 Generalization to Real-World Traffic

Although the NSL-KDD dataset provides a valuable testbed, it does not capture the dynamic nature of real-world network traffic. The static nature of the data limits the generalizability of findings. Future iterations of this research must evaluate the hybrid IDS against live traffic, accounting for encrypted payloads, polymorphic malware, and traffic variability.

6.4 Implications for Deployment

Given its relatively high detection accuracy and low FPR, the hybrid model could serve as a second-layer IDS in layered security architectures. For example:

- **Perimeter defense**: Initial packet filtering using signature rules
- **Behavioral detection**: Hybrid ML system scans flagged flows
- **Forensics**: Alerts fed into SIEM tools for correlation and response

Overall, the model demonstrates practical viability, but real-world implementation will require careful integration with monitoring, alerting, and incident response systems.

---

7. Conclusion

This study presents a hybrid intrusion detection system (IDS) that integrates two widely used machine learning algorithms—Random Forest and k-Nearest Neighbors (k-NN)—to address the limitations of traditional signature-based and anomaly-based IDS approaches. Through comprehensive experimentation using the NSL-KDD dataset, we demonstrate that the hybrid model not only improves detection accuracy (achieving 95.4%) but also significantly reduces false-positive rates (to 2.7%) compared to its constituent algorithms. These findings validate the hypotheses that ensemble methods, particularly those combining global and local learners, can provide a more balanced and reliable solution for network security monitoring.

The hybrid approach outperformed individual models in both overall metrics and category-specific evaluations, particularly excelling in identifying less frequent but highly damaging attacks like Remote-to-Local (R2L) and User-to-Root (U2R). This adaptability is essential in modern cybersecurity environments where attackers often deploy evasive, low-frequency threats to bypass conventional IDS filters. Furthermore, the consistent performance across multiple attack types underscores the model's robustness and generalization capability.

From a practical perspective, the reduction in false positives translates directly to reduced operational overhead for security analysts and lower likelihood of alert fatigue—a significant advantage in high-traffic enterprise networks. The ability to integrate such a model with packet inspection tools and log management systems (e.g., SIEM platforms) offers a path toward real-time deployment in security operations centers.

However, it is important to acknowledge the limitations of this research. The use of the NSL-KDD dataset, while convenient and standardized, does not fully capture the complexity of real-world network traffic. Future research should evaluate the hybrid model using live traffic datasets or simulated environments that incorporate modern protocols, encrypted payloads, and polymorphic attacks. Additionally, the computational cost of ensemble learning—particularly from k-NN's instance-based querying—requires optimization before production-level deployment.

Further enhancements may include:

- Incorporation of deep learning models such as autoencoders or LSTMs for temporal pattern recognition.
- Use of adaptive learning or online training to evolve the model in real-time based on newly detected threats.
- Integration with threat intelligence feeds and external anomaly scoring systems to enhance contextual awareness.

In conclusion, this research contributes a scalable and adaptable machine learning-based IDS framework suitable for modern enterprise networks. By combining the interpretability and strength of Random Forest with the instance-based sensitivity of k-NN, the hybrid system offers a robust alternative to traditional IDS models. Its demonstrated performance paves the way for future implementations that can evolve with the rapidly changing cybersecurity threat landscape, ultimately leading to more secure, intelligent, and proactive defense systems.

References

1. Anderson, B., & McGrew, D. (2017). Machine learning for encrypted malware traffic classification: Accounting for noisy labels and non-stationarity. *Proceedings of the ACM Workshop on Artificial Intelligence and Security*, 81–87. https://doi.org/10.1145/3128572.3140444
2. Munnangi, S. (2016). Adaptive case management (ACM) revolution. NeuroQuantology, 14(4), 844–850. https://doi.org/10.48047/nq.2016.14.4.974
3. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
4. 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. https://doi.org/10.1109/COMST.2015.2494502
5. Bellamkonda, S. (2017). Optimizing Your Network: A Deep Dive into Switches. NeuroQuantology, 15(1), 129-133.
6. Cover, T. M., & Hart, P. E. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1), 21–27. https://doi.org/10.1109/TIT.1967.1053964
7. Lippmann, R. P., Haines, J. W., Fried, D. J., Korba, J., & Das, K. (2000). The 1999 DARPA off-line intrusion detection evaluation. *Computer Networks*, 34(4), 579–595. https://doi.org/10.1016/S1389-1286(00)00139-0
8. Mukkamala, S., Sung, A. H., & Abraham, A. (2005). Intrusion detection using ensemble of soft computing paradigms. In *Proceedings of the 2005 International Conference on Intelligent Systems Design and Applications* (pp. 239–248). https://doi.org/10.1109/ISDA.2005.85

9. Patcha, A., & Park, J. M. (2007). An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12), 3448–3470. https://doi.org/10.1016/j.comnet.2007.02.001

10. Shon, T., & Moon, J. (2007). A hybrid machine learning approach to network anomaly detection. *Information Sciences*, 177(18), 3799–3821. https://doi.org/10.1016/j.ins.2007.03.025

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

12. Lee, W., & Stolfo, S. J. (2000). A framework for constructing features and models for intrusion detection systems. *ACM Transactions on Information and System Security*, 3(4), 227–261. https://doi.org/10.1145/382912.382914

13. Lazarevic, A., Ertoz, L., Kumar, V., Ozgur, A., & Srivastava, J. (2003). A comparative study of anomaly detection schemes in network intrusion detection. In *Proceedings of the 2003 SIAM International Conference on Data Mining* (pp. 25–36). https://doi.org/10.1137/1.9781611972733.3

14. Peddabachigari, S., Abraham, A., Grosan, C., & Thomas, J. (2007). Modeling intrusion detection system using hybrid intelligent systems. *Journal of Network and Computer Applications*, 30(1), 114–132. https://doi.org/10.1016/j.jnca.2005.09.001

15. Debar, H., Dacier, M., & Wespi, A. (2000). A revised taxonomy for intrusion-detection systems. *Annales des Télécommunications*, 55(7–8), 361–378. https://doi.org/10.1007/BF03192726