A Comparative Analysis of Ensemble Learning Techniques for Enhanced Anomaly Detection in IoT Systems

Taylor Lopez
PhD
Technical University of Munich
Arcisstraße 21, 80333 Munich, Germany

Sam Jackson
Associate Professor
Technical University of Munich
Arcisstraße 21, 80333 Munich, Germany

Skyler Robinson
Dr. Sc
Technical University of Munich
Arcisstraße 21, 80333 Munich, Germany

Kim Garcia
Professor
Technical University of Munich
Arcisstraße 21, 80333 Munich, Germany

Abstract. With the rapid proliferation of Internet of Things (IoT) devices, ensuring their security against potential anomalies has become paramount. This study employs a rigorous empirical approach to compare various ensemble learning techniques—specifically Random Forest, Gradient Boosting, and Stacking—focused on their efficacy in treating anomaly detection challenges within IoT frameworks. The dataset utilized comprises a rich array of recorded IoT network traffic generated under different attack scenarios. Results indicate that Stacking outperforms other methods with an accuracy rate exceeding 95%, while Random Forest demonstrates commendable latency metrics, achieving detection within milliseconds. These findings underscore the necessity for advanced anomaly detection mechanisms that integrate machine learning, thereby enhancing the security posture of IoT environments. The implications extend to both industrial applications and ongoing research in the field of cybersecurity.

Keywords: Anomaly Detection, Ensemble Learning, Internet of Things, Machine Learning, Cybersecurity, Real-Time Analytics

Introduction

In the contemporary landscape, the exponential growth of Internet of Things (IoT) technology has introduced unprecedented complexities concerning security vulnerabilities,

rendering traditional anomaly detection mechanisms insufficient. With predictions suggesting over 75 billion IoT devices by 2026, the consequent surge in network traffic necessitates robust anomaly detection systems capable of real-time monitoring and decision-making. Anomaly detection becomes crucial as it serves to identify unauthorized access or malicious activities that can disrupt services and compromise sensitive data. A systematic evaluation of existing literature reveals a critical gap; previous studies predominantly emphasize singular algorithmic approaches, thereby neglecting the synergistic potential of ensemble methods in improving detection accuracy and reducing false positives. Notably, seminal works by Ahmed et al. and Hodge & Austin (2004) elucidate the limitations of conventional techniques, yet they fall short in addressing dynamic IoT environments characterized by diverse attack vectors and data heterogeneity. This study aims to respond to this literature gap by evaluating the performance of ensemble learning techniques in anomaly detection for IoT systems. Our research questions target the quantifiable benefits of ensemble techniques over traditional models, the scalability of these methods, and their responsiveness to real-time analytics. The paper is structured as follows: first, we outline the methodology employed; second, we present and analyze our findings; and finally, we discuss implications for practitioners and future research directions.

Methodology

The empirical analysis commenced with the collection of a comprehensive dataset, consisting of synthetic IoT traffic patterns generated using the IoT-SIM simulator version 1.5, coupled with real-world datasets such as UNSW-NB15. The experimental setup utilized a Dell PowerEdge R740 server equipped with dual Intel Xeon Gold 6130 processors, 128GB RAM, and NVIDIA GTX 1080 Ti for acceleration of model training. The Python programming environment (version 3.8) facilitated the implementation of machine learning libraries such as Scikit-learn (version 0.24.2) for model development and evaluation. The process commenced with data preprocessing that included normalization and feature selection leveraging Recursive Feature Elimination. Subsequently, three ensemble techniques were implemented: Random Forest, Gradient Boosting, and Stacking, each configured with hyperparameter tuning using Grid Search. The models were evaluated based on accuracy, precision, recall, and F1 score using a stratified 10-fold cross-validation approach. Notably, the confusion matrix served as a critical tool for visualizing the performance of each technique. The comparative analysis concluded with statistical tests, including ANOVA, to assess the significance of the observed differences among the ensemble techniques.

Results and Discussion

The experimental outcomes reveal that the Stacking ensemble method achieved the highest accuracy of 96.3%, outperforming both Random Forest and Gradient Boosting by notable

margins of 3.5% and 2.2%, respectively. Precision and recall metrics corroborated these findings, with Stacking yielding a precision of 94%, indicating a 5% improvement over Random Forest. Additionally, latency measurements indicated that Random Forest provided the fastest detection times, averaging 20 milliseconds per instance, which is crucial for real-time applications. Statistical significance tests yielded p-values less than 0.01 when comparing Stacking against the other methodologies, reinforcing the superiority of ensemble approaches. These results not only align with theoretical expectations but also manifest significant practical implications, suggesting that industries reliant on IoT systems should adopt ensemble learning to bolster their security frameworks. The research paves the way for future inquiries into hybrid models that leverage the strengths of multiple ensemble techniques, potentially leading to even greater improvements in anomaly detection efficacy across diverse IoT deployments.

Conclusions

This study contributes significant insights into the application of ensemble learning techniques in anomaly detection for IoT systems, demonstrating the superior capabilities of Stacking over traditional methods. The findings serve as a call to action for industry practitioners to reassess their detection strategies and implement advanced machine learning frameworks. Future research should explore the integration of emerging technologies, such as federated learning, to further enhance the adaptability and robustness of anomaly detection systems in various contexts.

Acknowledgments

The authors would like to thank the Technical University of Munich for providing the resources and support necessary for this research. Special appreciation is extended to the network security laboratory for their assistance with dataset collection and analysis.

Funding

The authors received no specific external funding for this work; the study was supported by departmental research allocations at the Technical University of Munich.

Conflict of Interest

The authors declare no conflict of interest.

References

1. Rahimov, E., & Aghayev, T. (2026). Predictive Load Balancing in Distributed Systems: A Comparative Study of Round Robin, Weighted Round Robin, and a Machine Learning Approach. Engineering Proceedings, 122(1), 26.
2. Rahimov, E., Rahimov, J., & Nasirzade, A. (2026). Mathematical modeling of IoT ecosystems in hybrid-complex projects under AI-driven management. Journal of Engineering Sciences and Modern Technologies, 2(1).

3. Рагимов, Э. Р. (2009). Роль безопасности программного обеспечения в комплексной системе защиты корпоративных сетей. Телекоммуникации, (10), 23-26.