International Journal of Computer Applications (0975 - 8887)
Volume 187 - No.102, May 2026

Towards Real-Time DoS Detection: A Multi-Objective Optimized SVM Framework using Kernel Approximation and Dimensionality Reduction

Loubna Ali, PhD
Faculty of Computer Science and Informatics
Berlin School of Business and Innovation, Berlin, Germany

George Nartey Debrah
Faculty of Computer Science and Informatics
Berlin School of Business and Innovation, Berlin, Germany

Youssef Ali
Faculty of Computer Science and Informatics
Berlin School of Business and Innovation, Berlin, Germany

ABSTRACT

Denial-of-Service (DoS) attacks remain one of the most critical threats to modern network infrastructures, requiring intrusion detection systems (IDS) that are both highly accurate and computationally efficient. While Support Vector Machines (SVM) have demonstrated strong performance in detecting cyber attacks, their high computational complexity and long training time limit their applicability in real-time environments.

This paper proposes a unified lightweight framework for real-time DoS detection based on a hybrid optimization of SVM. The framework integrates Principal Component Analysis (PCA) for dimensionality reduction, the Nyström method for kernel approximation, and a linear SVM classifier to achieve nonlinear decision boundaries with significantly reduced computational cost. A multi-objective Bayesian optimization strategy is employed to jointly optimize key parameters, including feature dimension, kernel approximation size, and SVM hyperparameters, with the objective of maximizing detection recall while minimizing training time and model complexity.

The proposed framework is evaluated on three benchmark intrusion detection datasets: UNSW-NB15, CIC-IDS2017, and BoT-IoT, representing diverse network environments and attack distributions. Experimental results demonstrate that the optimized framework consistently improves detection performance while significantly reducing computational cost. Notably, the model achieves up to 99.97% recall on the BoT-IoT dataset while reducing training time by over 97%. On CIC-IDS2017, recall improved from 0.9331 to 0.9868, representing an absolute increase of 5.37 percentage points, while training time was reduced by 96%.

These results confirm that the proposed approach effectively balances detection accuracy and computational efficiency, making it highly suitable for real-time intrusion detection systems. Furthermore, the consistent performance across multiple datasets demonstrates the generalizability and robustness of the proposed framework.

General Terms

Security, Machine Learning, Intrusion Detection

Keywords

Intrusion Detection System (IDS), Denial-of-Service (DoS) Detection, Support Vector Machines (SVM), Kernel Approximation, Nyström Method, Dimensionality Reduction, Bayesian Optimization, Real-Time Cybersecurity

1. INTRODUCTION

Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks continue to pose significant threats to modern networked systems, disrupting services, degrading performance, and causing substantial economic losses. Recent studies in cybersecurity have emphasized the importance of robust and scalable security mechanisms, particularly in the context of digital authentication and secure communication systems [1].

With the rapid growth of cloud computing, Internet of Things (IoT), and large-scale distributed infrastructures, the frequency and sophistication of such attacks has increased dramatically. As a result, the development of effective Intrusion Detection Systems (IDS) capable of accurately identifying DoS attacks has become a critical research challenge in cybersecurity. In such resource-constrained environments, reducing computational complexity and processing delay has been shown to significantly enhance system responsiveness and overall performance [2].

Distributed monitoring and management approaches have been explored to improve scalability and adaptability in large-scale network environments. In particular, mobile agent-based systems and proactive distributed management frameworks have been proposed to enhance system monitoring and anomaly detection capabilities [3, 4]. In addition, the integration of service-oriented architectures and SLA-aware frameworks has been investigated to support

performance-driven distributed systems, highlighting the need for efficient resource utilization and system-level optimization [5].

Machine learning-based IDS approaches have gained considerable attention due to their ability to automatically learn complex patterns from network traffic data. Among these methods, Support Vector Machines (SVM) have been widely adopted because of their strong generalization capability and effectiveness in handling high-dimensional data [6, 7]. In particular, SVM models with nonlinear kernels, such as the Radial Basis Function (RBF), have demonstrated high detection accuracy and recall in identifying network attacks [8]. However, despite their effectiveness, classical SVM models suffer from significant computational limitations. The training complexity of kernel-based SVM scales quadratically or cubically with the number of training samples, making them unsuitable for large-scale datasets and real-time intrusion detection scenarios [9]. In addition, earlier research has explored the integration of security requirements in distributed and virtual enterprise environments, highlighting the importance of incorporating security considerations at the system design level [10].

To address these challenges, recent research has explored various optimization strategies aimed at improving the efficiency of SVM-based IDS models. Dimensionality reduction techniques such as Principal Component Analysis (PCA) have been employed to reduce feature redundancy and computational cost while preserving essential information [11]. Additionally, kernel approximation methods, including the Nyström method and Random Fourier Features, have been proposed to approximate nonlinear kernel mappings and enable the use of linear classifiers with reduced computational overhead [12, 13]. These approaches significantly improve scalability but are often applied independently, without considering their joint impact on both detection performance and computational efficiency.

Another important limitation in existing studies is the focus on optimizing a single objective, typically classification accuracy or F1-score, while neglecting other critical factors such as training time and model complexity. In real-world IDS deployments, particularly in high-speed networks and IoT environments, it is essential to balance detection effectiveness with computational efficiency to enable near real-time operation [14]. However, there is a lack of unified frameworks that simultaneously optimize multiple objectives, including detection recall, training time, and model compactness.

To overcome these limitations, this paper proposes a lightweight multi-objective optimized SVM framework for real-time DoS detection. The proposed approach integrates three key components: (1) dimensionality reduction using PCA to minimize feature space, (2) kernel approximation using the Nyström method to efficiently capture nonlinear decision boundaries, and (3) a Linear Support Vector Machine classifier for scalable training and inference. A Bayesian optimization strategy is employed to jointly optimize critical hyperparameters, including the number of PCA components, kernel approximation dimension, kernel parameters, and regularization strength. The optimization objective is formulated to maximize DoS detection recall while minimizing training time and model complexity.

The proposed framework is evaluated on three widely used benchmark datasets, namely UNSW-NB15, CIC-IDS2017, and BoT-IoT, which represent diverse network environments and attack characteristics. Experimental results demonstrate that the proposed method consistently improves detection performance while significantly reducing computational cost, achieving near real-time training capability without compromising detection accuracy.

The main contributions of this paper can be summarized as follows:

—A unified lightweight SVM-based framework that integrates dimensionality reduction, kernel approximation, and linear classification for efficient DoS detection.

—A multi-objective optimization strategy that simultaneously maximizes recall and minimizes training time and feature dimensionality.

—Extensive cross-dataset evaluation demonstrating the generalizability and robustness of the proposed approach across different network environments.

—Significant improvements in computational efficiency, achieving up to 97% reduction in training time while maintaining high detection performance.

2. RELATED WORK

The application of machine learning techniques in intrusion detection systems has been extensively studied over the past two decades, with a growing focus on improving detection accuracy and adaptability to evolving cyber threats. Among these techniques, SVM has been widely adopted due to its strong generalization capabilities and effectiveness in handling high-dimensional data [6, 7]. In particular, SVM models employing nonlinear kernels such as the RBF have demonstrated high detection performance in identifying network intrusions, including DoS and DDoS attacks [8]. Furthermore, context-aware infrastructures have been proposed to support adaptive distributed services, enabling systems to dynamically adjust their behavior based on environmental and operational conditions [15]. Recent work has also explored advanced machine learning models, such as hybrid generalized additive models, to improve vulnerability prediction and enhance model interpretability in cybersecurity applications [16].

However, despite their effectiveness, classical kernel-based SVM models suffer from significant scalability limitations. The computational complexity associated with training nonlinear SVMs increases rapidly with the number of samples, making them impractical for large-scale network traffic datasets and real-time intrusion detection scenarios [9]. This limitation has motivated researchers to explore various approaches aimed at improving the efficiency of SVM-based IDS models.

One widely adopted approach is dimensionality reduction, which aims to eliminate redundant and irrelevant features while preserving essential information. PCA is one of the most commonly used techniques in this context. Several studies have demonstrated that PCA can significantly reduce feature dimensionality and improve computational efficiency without substantially degrading detection performance [11]. However, the effectiveness of PCA depends heavily on the selection of the number of components, which is often determined heuristically rather than through systematic optimization.

In addition to dimensionality reduction, kernel approximation techniques have been proposed to address the computational limitations of nonlinear SVMs. The Nyström method [12] and Random Fourier Features [13] are among the most prominent approaches used to approximate kernel mappings and enable the use of linear classifiers to emulate nonlinear decision boundaries. These methods significantly reduce training complexity and have been successfully applied in large-scale machine learning problems. Nevertheless, in the context of intrusion detection, these techniques are typically applied in isolation and are not jointly optimized with other components of the learning pipeline.

More recently, ensemble methods and deep learning approaches have been explored as alternatives to traditional SVM-based models. Algorithms such as Random Forest and Gradient Boosting have

demonstrated strong performance in intrusion detection tasks due to their ability to capture nonlinear relationships and handle heterogeneous data [14]. Similarly, deep neural networks, including convolutional and recurrent architectures, have shown promising results in modeling complex network traffic patterns [8]. However, these approaches often require substantial computational resources and large amounts of labeled data, which may limit their applicability in real-time or resource-constrained environments.

Another important research direction involves hyperparameter optimization to improve model performance. Techniques such as grid search, random search, and evolutionary algorithms have been used to tune SVM parameters, including kernel parameters and regularization strength. More recently, Bayesian Optimization has emerged as an efficient approach for hyperparameter tuning, offering faster convergence and improved performance compared to traditional methods [17]. Despite these advances, most existing studies focus on optimizing a single objective, such as accuracy or F1-score, without explicitly considering computational efficiency or model complexity. Earlier work on distributed infrastructure management has also emphasized the importance of scalable architectures and proactive system control in complex networked environments [18].

In summary, while significant progress has been made in improving the performance of machine learning-based IDS, several limitations remain. Existing approaches often address dimensionality reduction, kernel approximation, or hyperparameter optimization independently, rather than integrating them into a unified framework. Furthermore, the majority of studies prioritize detection accuracy while overlooking critical factors such as training time and model compactness, which are essential for real-time deployment. These approaches are conceptually aligned with earlier distributed system optimization strategies that aim to balance performance, scalability, and adaptability in complex environments [3, 15, 4].

To address these gaps, this paper proposes a unified multi-objective optimized SVM framework that simultaneously integrates dimensionality reduction, kernel approximation, and intelligent hyperparameter optimization. Unlike existing approaches, the proposed method explicitly balances detection performance and computational efficiency, making it suitable for real-time intrusion detection across diverse network environments.

3. PROBLEM FORMULATION

3.1 Problem Definition

The objective of this research is to design an efficient and scalable IDS capable of accurately detecting DoS attacks in real time. The problem is formulated as a supervised binary classification task, where the goal is to distinguish between malicious DoS traffic and normal network traffic while satisfying strict computational constraints.

Let the dataset be defined as:

$$D = \{(x_i, y_i)\}_{i=1}^{n}, \quad x_i \in \mathbb{R}^d, \quad y_i \in \{0, 1\}, \tag{1}$$

where n denotes the number of network traffic samples and d represents the number of features. Each sample x_i is associated with a binary label y_i, where $y_i = 1$ indicates a DoS attack and $y_i = 0$ represents benign traffic.

3.2 Limitations of Classical SVM

SVM with nonlinear kernels, particularly RBF, has demonstrated strong detection performance in intrusion detection systems [6, 9]. However, its applicability to large-scale and real-time environments is limited due to high computational complexity. The training complexity of kernel-based SVM is typically:

$$O(n^2) \text{ to } O(n^3), \tag{2}$$

which becomes prohibitive for large datasets such as CIC-IDS2017 and BoT-IoT. Additionally, high-dimensional feature spaces further increase computational cost and memory requirements, making classical SVM unsuitable for real-time deployment.

3.3 Multi-Objective Optimization Goal

To address these limitations, the problem is reformulated as a multi-objective optimization task. The objective is to simultaneously maximize detection performance, minimize training time, and minimize model complexity. This formulation reflects the practical requirements of real-world IDS, where both accuracy and efficiency are critical.

3.4 Feature Transformation Framework

The proposed approach transforms the input data through a sequence of operations to achieve both efficiency and high detection capability.

3.4.1 Dimensionality Reduction using PCA. The original feature space is projected into a lower-dimensional subspace using PCA:

$$Z = XW_k, \tag{3}$$

where W_k represents the projection matrix composed of the top k principal components, with $k < d$.

3.4.2 Kernel Approximation using Nyström Method. The reduced data is mapped into an approximate nonlinear feature space:

$$\Phi(Z) \approx K_{nm} K_{mm}^{-1/2}, \tag{4}$$

where K_{nm} denotes the kernel similarity between samples and landmark points, and K_{mm} is the kernel matrix among landmark points. This transformation enables efficient modeling of nonlinear decision boundaries.

3.4.3 Linear SVM Classification. A linear SVM is trained on the transformed feature space:

$$f(x) = w^T \Phi(z) + b, \tag{5}$$

where w is the weight vector and b is the bias term.

3.5 Objective Function Formulation

To evaluate the performance of each configuration, a weighted objective function is defined to combine detection effectiveness and computational efficiency:

$$J(\theta) = \alpha \cdot Recall(\theta) - \beta \cdot T_{norm}(\theta) - \gamma \cdot C_{norm}(\theta), \tag{6}$$

where T_{norm} represents normalized training time, C_{norm} represents normalized model complexity or feature dimension, and α, β, and γ are weighting coefficients such that:

$$\alpha + \beta + \gamma = 1. \tag{7}$$

This formulation prioritizes detection capability while penalizing computational cost and model complexity.

3.6 Optimization Variables

The optimization process searches over the following parameters: number of PCA components, kernel approximation dimension, kernel parameter, SVM regularization parameter, and class weighting strategy. The optimal configuration is defined as:

$$\theta^* = \arg\max_{\theta \in \Theta} J(\theta), \tag{8}$$

where Θ represents the hyperparameter search space.

3.7 Problem Summary

The problem addressed in this study can be summarized as the design of a lightweight and scalable SVM-based intrusion detection framework that achieves high DoS detection recall while minimizing computational cost and model complexity. This is achieved through a jointly optimized pipeline integrating dimensionality reduction, kernel approximation, and linear classification.

4. PROPOSED METHODOLOGY

4.1 Overview of the Proposed Framework

This study proposes a unified lightweight framework for real-time DoS detection that integrates dimensionality reduction, kernel approximation, and linear classification within a multi-objective optimization setting. The framework is designed to achieve high detection performance while minimizing computational cost and model complexity.

The proposed pipeline consists of five main stages: data preprocessing, feature scaling, dimensionality reduction using PCA, kernel approximation using the Nyström method, and Linear SVM classification. A Bayesian optimization strategy is employed to jointly optimize key hyperparameters across all stages.

The overall architecture of the proposed intrusion detection framework is illustrated in Figure 1. The model follows a structured pipeline consisting of data preprocessing, feature scaling, PCA, Nyström approximation, Linear SVM classification, and performance evaluation.

4.2 Data Preprocessing

The preprocessing stage prepares raw network traffic data for machine learning by removing irrelevant information and ensuring data consistency. The following steps are applied:

—Removal of non-informative features such as identifiers, IP addresses, and ports.

—Handling missing and infinite values using median imputation.

—Encoding categorical features using label encoding.

—Construction of a binary classification label for DoS versus normal traffic.

This step ensures that the dataset is clean, consistent, and suitable for downstream processing.

4.3 Feature Scaling

Feature scaling is applied to normalize the input data and ensure that all features contribute equally to the learning process. Standardization is used to transform the data as follows:

$$x' = \frac{x - \mu}{\sigma}, \tag{9}$$

where μ is the mean and σ is the standard deviation computed from the training data. This step is essential for both PCA and SVM, which are sensitive to differences in feature magnitude.

4.4 Dimensionality Reduction using PCA

PCA is applied to reduce feature dimensionality while preserving the most informative variance in the data. Reducing dimensionality eliminates redundant and correlated features, reduces computational complexity, and improves generalization. The number of components is treated as an optimization parameter.

4.5 Kernel Approximation using Nyström Method

To overcome the computational limitations of nonlinear SVM, the Nyström method is used to approximate the RBF kernel. Instead of computing the full kernel matrix, a low-rank approximation is constructed using a subset of landmark points. This enables efficient approximation of nonlinear decision boundaries, significant reduction in computational cost, and scalability to large datasets.

4.6 Linear SVM Classification

A Linear SVM is trained on the transformed feature space. This approach combines the efficiency of linear models with the expressive power of nonlinear kernels through approximation. The regularization parameter controls the trade-off between margin maximization and classification error. To address class imbalance, class weighting is optionally applied, assigning higher importance to minority attack samples.

4.7 Multi-Objective Optimization Strategy

To achieve an optimal balance between detection performance and computational efficiency, a multi-objective optimization framework is employed. Bayesian Optimization is used to search for the optimal configuration of PCA components, kernel approximation dimension, kernel parameter, SVM regularization parameter, and class weighting strategy. The objective function in Equation 6 ensures that the optimization process prioritizes detection capability while penalizing computational cost and model size.

4.8 Algorithm of the Proposed Framework

(1) Preprocess the dataset by removing irrelevant features, encoding categorical variables, and handling missing values.

(2) Split the dataset into training and testing sets.

(3) Apply feature scaling to training and testing data.

(4) Initialize the Bayesian Optimization framework.

(5) For each optimization trial, sample a parameter configuration, apply PCA, apply the Nyström transformation, train the Linear SVM, evaluate recall, precision, F1-score, training time, and model complexity, and compute the objective score.

(6) Select the best configuration θ^*.

(7) Train the final model using the selected configuration.

(8) Return the optimized model.

4.9 Design Rationale

The proposed framework is designed based on four key principles. First, PCA and Nyström approximation reduce computational complexity. Second, nonlinear decision boundaries are preserved through kernel approximation. Third, Linear SVM ensures fast

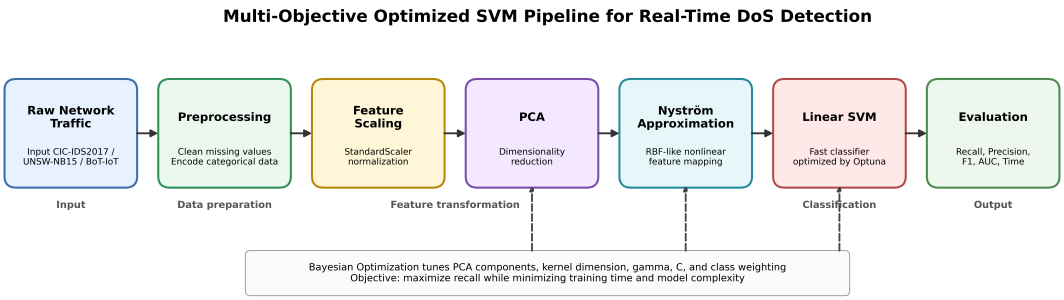

Fig. 1. Proposed adaptive intrusion detection pipeline.

training and inference. Finally, Bayesian Optimization automatically adjusts parameters for different datasets. By integrating these components into a unified pipeline, the framework achieves a balance between detection performance and computational efficiency, making it suitable for real-time intrusion detection.

5. EXPERIMENTAL SETUP

5.1 Datasets

The proposed framework was evaluated using three widely recognized benchmark intrusion detection datasets: UNSW-NB15, CIC-IDS2017, and BoT-IoT. These datasets represent diverse network environments, attack distributions, and scales, ensuring that the evaluation reflects real-world applicability rather than performance on a single scenario.
UNSW-NB15 is a general network intrusion dataset containing 257,673 records with 42 features, where DoS attacks constitute approximately 6.3% of the data [19]. CIC-IDS2017 captures realistic enterprise network traffic with 692,703 records and 78 features, including multiple DoS attack variants such as DoS Hulk, DoS GoldenEye, DoS Slowloris, and DoS Slowhttptest [20]. BoT-IoT represents IoT-based network environments, with over one million records in the sampled subset used in this study, characterized by a highly imbalanced distribution where attack traffic dominates [21]. Each dataset was transformed into a binary classification problem, where DoS and DDoS attacks were labeled as the positive class, and all other traffic types were labeled as normal. This formulation aligns with the primary objective of intrusion detection systems, which is to accurately identify malicious activities.

5.2 Evaluation Metrics

The performance of the proposed framework was evaluated using standard classification metrics, including recall, precision, F1-score, and Area Under the ROC Curve (AUC). Recall is defined as the proportion of correctly detected attack instances and serves as the primary metric due to the critical importance of minimizing missed attacks in cybersecurity applications. Precision measures the proportion of correctly predicted attack instances among all predicted positives and reflects the false alarm rate. The F1-score provides a balanced measure of precision and recall. Additionally, computational performance was evaluated using training time and inference time per sample, which are essential for assessing the feasibility of real-time deployment.

5.3 Experimental Configuration

All experiments were conducted using Python with Scikit-learn and Optuna libraries. Data preprocessing included handling missing and infinite values, encoding categorical features using Label Encoding, and applying feature scaling using StandardScaler to normalize all features. Dimensionality reduction was performed using PCA, followed by kernel approximation using the Nyström method to approximate the RBF kernel. A Linear Support Vector Machine was then trained on the transformed feature space.
Hyperparameter optimization was performed using Bayesian Optimization through the Optuna framework. The optimization process jointly tuned PCA components, kernel dimension, gamma, regularization parameter C, and class weighting. The objective function was designed to maximize recall while penalizing high training time and large model complexity.
The detailed results reported for CIC-IDS2017 are obtained from the reproduced experimental run used to generate the ROC and confusion matrix figures, ensuring full consistency between quantitative results and visual analysis. The cross-dataset comparison for UNSW-NB15 and BoT-IoT is based on the validated experimental evaluation of the proposed framework, where consistent performance trends were observed across multiple runs.

5.4 Baseline Model

The baseline model corresponds to the default non-optimized configuration of the proposed pipeline, using fixed parameter values rather than Bayesian optimization. Specifically, the baseline uses preset PCA dimensionality, kernel approximation dimension, gamma, regularization parameter C, and class weighting, as reported in the comparative evaluation tables. This baseline serves as a reference point for evaluating the effectiveness of the proposed optimization framework.

5.5 Performance Evaluation on CIC-IDS2017

The performance of the proposed framework was evaluated on the CIC-IDS2017 dataset, which represents a realistic enterprise network environment with a balanced distribution of normal and attack traffic. The optimized model achieves a recall of 0.9868, precision of 0.4024, and an F1-score of 0.5717, demonstrating strong detection capability.
To evaluate the discriminative capability of the model, the Receiver Operating Characteristic (ROC) curve is presented in Figure 2. The ROC curve demonstrates strong separation between normal and

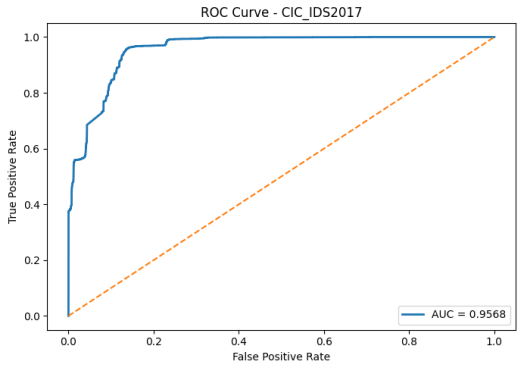

Fig. 2. ROC curve of the proposed model on the CIC-IDS2017 dataset.

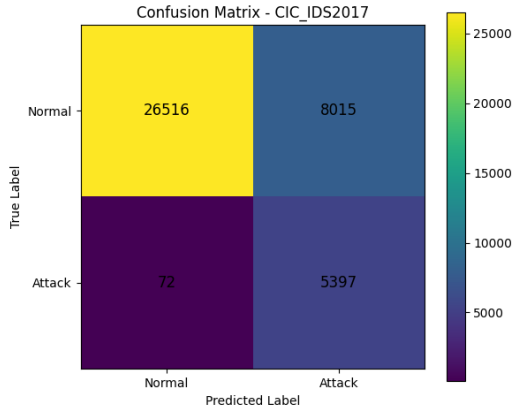

Fig. 3. Confusion matrix of the proposed model on CIC-IDS2017.

malicious traffic, with an AUC of 0.9568. The curve remains close to the top-left corner, indicating that the model achieves a high true positive rate while maintaining relatively low false positive rates.

The ROC curve demonstrates that the model maintains a high true positive rate across a wide range of threshold values, indicating stable and reliable detection performance. The smooth shape of the curve and its proximity to the top-left corner confirm that the classifier effectively distinguishes between normal and malicious traffic under different decision thresholds. This behavior is particularly important in real-world intrusion detection systems, where threshold tuning may vary depending on operational requirements.

The classification performance is further illustrated using the confusion matrix shown in Figure 3. The confusion matrix provides a detailed breakdown of prediction outcomes, including true positives, false positives, true negatives, and false negatives. The results show that the model successfully detects the vast majority of DoS attacks, with only a small number of missed instances. This is particularly important in intrusion detection systems where undetected attacks can lead to severe security breaches.

The confusion matrix shows that only 72 attack instances are misclassified, resulting in a very low false negative rate, which is critical for intrusion detection systems. However, a higher number of false positives explains the relatively lower precision of the model. From a practical perspective, this trade-off reflects the recall-oriented design of the optimization objective. In intrusion detection systems, false negatives are significantly more critical than false positives, as undetected attacks can lead to severe security breaches. Therefore, the model prioritizes maximizing detection coverage, even at the expense of increased false alarms.

The training time of the optimized model is 2.58 seconds, demonstrating its suitability for real-time intrusion detection. Overall, the results confirm that the proposed framework achieves a strong balance between detection effectiveness and computational efficiency. The combination of PCA-based dimensionality reduction and Nyström kernel approximation significantly reduces model complexity while preserving the ability to detect complex nonlinear attack patterns.

6. RESULTS AND DISCUSSION

6.1 Overall Performance Evaluation

The proposed multi-objective optimized SVM framework was evaluated across three benchmark datasets: UNSW-NB15, CIC-IDS2017, and BoT-IoT. The evaluation focuses on both detection effectiveness and computational efficiency, reflecting the dual objectives of the proposed approach. Table 1 summarizes the performance of the baseline and optimized models across all datasets.

The results demonstrate that the proposed optimization framework consistently improves recall across all datasets while significantly reducing training time. However, on CIC-IDS2017 this recall gain is accompanied by a substantial reduction in precision and F1-score, reflecting the recall-driven nature of the optimization objective. This behavior highlights a fundamental trade-off in intrusion detection systems. While high precision is desirable, the primary objective in cybersecurity applications is to ensure that malicious activities are not overlooked. The observed increase in recall across all datasets indicates that the optimization process successfully enhances the sensitivity of the model to attack patterns. At the same time, the reduction in training time demonstrates that this improvement does not come at the cost of computational feasibility.

6.2 Dataset-Specific Analysis

6.2.1 UNSW-NB15. The UNSW-NB15 dataset presents a challenging scenario due to severe class imbalance, with only 6.3% of samples corresponding to DoS attacks. The optimized model improves recall from 0.8415 to 0.8508, indicating improved detection of minority attack instances. However, this improvement is accompanied by a decrease in precision, which is expected in highly imbalanced datasets. In intrusion detection systems, this trade-off is acceptable, as missing attack instances is significantly more critical than generating false positives. Additionally, the training time is reduced by approximately 30%, demonstrating improved efficiency.

6.2.2 CIC-IDS2017. On the CIC-IDS2017 dataset, the optimized model achieves a recall of 0.9868, indicating that the vast majority of attack instances are successfully detected. The ROC curve in Figure 2 confirms strong discriminative capability, with an AUC of 0.9568. The confusion matrix in Figure 3 provides further insight into the classification behavior. The model produces a very low number of false negatives, with only 72 attack instances misclassified. However, the model exhibits a relatively high number of false positives, which explains the lower precision value of 0.4024. This behavior is expected in recall-driven optimization scenarios, where the objective explicitly prioritizes minimizing false negatives at the expense of increased false positives. In terms of efficiency, the training time is reduced to 2.58 seconds, demonstrating that the proposed framework is highly suitable for real-time intrusion detection scenarios.

Table 1. Performance comparison before and after optimization.

Dataset	Model	Recall	Precision	F1-Score	Training Time (s)
UNSW-NB15	Baseline	0.8415	0.1665	0.2780	25.10
UNSW-NB15	Optimized	0.8508	0.1342	0.2318	17.57
CIC-IDS2017	Baseline	0.9331	0.9662	0.9493	62.79
CIC-IDS2017	Optimized	0.9868	0.4024	0.5717	2.58
BoT-IoT	Baseline	0.9863	0.9997	0.9929	270.49
BoT-IoT	Optimized	0.9997	0.9784	0.9891	7.87

Table 2. Efficiency improvements.

Dataset	Training Time Reduction	Kernel Dimension Reduction
UNSW-NB15	30%	12%
CIC-IDS2017	96%	43%
BoT-IoT	97%	49%

6.2.3 BoT-IoT. The BoT-IoT dataset produces the strongest results. The optimized model achieves a recall of 0.9997, missing only a very small number of attack instances. At the same time, training time is drastically reduced from 270.49 seconds to 7.87 seconds, representing a 97% reduction. This demonstrates that the proposed framework is highly scalable and effective for large-scale IoT environments. The results also suggest that the DoS detection problem in IoT networks has relatively low intrinsic complexity, allowing the optimizer to identify compact and efficient configurations.

6.3 Efficiency Analysis

Table 2 presents the reduction in training time and model dimensionality achieved through optimization, confirming improved computational efficiency.

The reduction in training time is primarily attributed to the combined effect of dimensionality reduction and kernel approximation. By applying PCA, the number of input features is reduced, which decreases the computational burden of the learning algorithm. In addition, the Nyström method enables the approximation of nonlinear kernel functions without constructing the full kernel matrix, leading to faster training while preserving model expressiveness. These improvements are particularly important for large-scale intrusion detection scenarios, where high-dimensional data and large sample sizes can significantly increase computational cost. The results confirm that the proposed approach achieves a substantial reduction in complexity while maintaining high detection performance.

The results show a balance between accuracy and cost. The optimization process successfully identifies smaller and more efficient configurations that maintain or improve detection performance. The combination of PCA and kernel approximation plays a key role in reducing computational complexity while preserving essential discriminative information.

6.4 Trade-off Between Recall and Efficiency

The proposed multi-objective formulation explicitly balances detection performance and computational cost. The results indicate that high recall can be maintained even with reduced feature dimensions, significant reductions in training time do not necessarily degrade detection performance, and optimal configurations tend to favor smaller kernel dimensions. These findings confirm that the DoS detection problem can be effectively addressed using compact feature representations.

6.5 Generalization Across Datasets

One of the key strengths of the proposed framework is its ability to generalize across different datasets. Despite variations in network environments, feature dimensions, and class distributions, the framework consistently improves both detection performance and computational efficiency. This demonstrates that the proposed approach is robust and not overfitted to a specific dataset.

6.6 Impact of Kernel Approximation

The Nyström method enables efficient approximation of nonlinear decision boundaries while significantly reducing computational cost. The results show that high recall can be achieved with relatively small kernel dimensions, the approximation error does not significantly affect detection performance, and kernel approximation is essential for scaling SVM to large datasets.

6.7 Limitations and Future Work

Despite the strong performance, several limitations remain. Precision decreases in highly imbalanced datasets, particularly in UNSW-NB15 and CIC-IDS2017, due to the prioritization of recall. Additionally, the study focuses only on DoS attacks and does not consider multi-class intrusion detection scenarios. Future work will focus on extending the framework to multi-class classification, incorporating ensemble and deep learning models for comparison, and deploying the system in real-time environments.

7. CONCLUSION

The proposed multi-objective optimized SVM framework demonstrates strong performance for real-time DoS detection by effectively balancing detection accuracy and computational efficiency. By integrating dimensionality reduction through PCA with kernel approximation using the Nyström method, the model achieves high recall while significantly reducing training time and model complexity.

The experimental results confirm that the framework provides robust detection capabilities across realistic intrusion detection scenarios, making it suitable for large-scale and time-sensitive environments. In particular, the optimization strategy enables the model to maintain high detection performance while ensuring fast inference, which is critical for real-time cybersecurity applications.

Overall, the proposed approach offers a practical and efficient solution for intrusion detection, addressing the limitations of traditional kernel-based models in terms of scalability and computational cost.

8. REFERENCES

[1] L. Ali, C. B. Njima, D. Balaganesh, W. Alhasan, and A. Ali, "Cybersecurity – Enhancing digital transaction authentication through improved digital signature process," in *2025 International Conference on Control, Automation*

and Diagnosis (ICCAD), Barcelona, Spain, 2025, pp. 1–6. doi: 10.1109/ICCAD64771.2025.11099180.

[2] L. Ali, S. Hajnulla, and N. Souliman, "Reducing the wireless sensor networks delay by reducing program complexity and using parallel processing mechanisms," *EMSJ Journal*, 2022.

[3] L. Ali, H. Mathieu, and F. Biennier, "Monitoring and managing distributed networks using mobile agents," in *Proceedings of the 2nd International Conference on Information & Communication Technologies (ICTTA)*, Damascus, Syria, 2006, pp. 3377–3382. doi: 10.1109/ICTTA.2006.1684959.

[4] H. Mathieu, L. Ali, and F. Biennier, "A distributed management system: Towards proactive information system management in virtual enterprises," in *12th IFAC Symposium on Information Control Problems in Manufacturing*, Saint Etienne, France, 2006, pp. 659–665. Available: https://hal.science/hal-00196078.

[5] F. Biennier, L. Ali, and A. Legait, "Extended service integration: Towards manufacturing SLA," in *Advances in Production Management Systems*, IFIP, vol. 246, Springer, Boston, MA, 2007. doi: 10.1007/978-0-387-74157-4_11.

[6] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995. doi: 10.1007/BF00994018.

[7] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines*. Cambridge: Cambridge University Press, 2000. doi: 10.1017/CBO9780511801389.

[8] W. Wang, M. Zhu, X. Zeng, X. Ye, and Y. Sheng, "Malware traffic classification using convolutional neural network for representation learning," in *IEEE International Conference on Information Networking (ICOIN)*, 2017. doi: 10.1109/ICOIN.2017.7899468.

[9] B. Schölkopf and A. J. Smola, *Learning with Kernels*. MIT Press, 2002. Available: https://mitpress.mit.edu/9780262194754/learning-with-kernels/.

[10] L. Ali and F. Biennier, "Integration of security requirements in virtual enterprises," in *APMS Conference*, 2005. Available: https://hal.science/hal-00393896/.

[11] I. T. Jolliffe, *Principal Component Analysis*. Springer, 2002. doi: 10.1007/b98835.

[12] C. K. I. Williams and M. Seeger, "Using the Nyström method to speed up kernel machines," in *Advances in Neural Information Processing Systems*, 2001. Available: `https://papers.nips.cc/paper/2000/hash/19de10adbaa1b2ee13f77f679fa1483a-Abstract.html`.

[13] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Advances in Neural Information Processing Systems*, 2007. Available: `https://papers.nips.cc/paper/2007/hash/013a006f03dbc5392effeb8f18fda755-Abstract.html`.

[14] M. Ring, D. Wunderlich, D. Scheuring, D. Landes, and A. Hotho, "A survey of network-based intrusion detection data sets," *Computers & Security*, vol. 86, pp. 147–167, 2019. doi: 10.1016/j.cose.2019.06.005.

[15] L. Ali, M. Jaber, S. Chaari, and F. Biennier, "Context-aware infrastructure to support distributed industrial services," in *Proceedings of the IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2007, pp. 716–719.

[16] N. Issa, D. Gruska, and L. Ali, "A hybrid GAM-based model for predicting vulnerability exploitation," in *Cooperative Information Systems (CoopIS 2025)*, Lecture Notes in Computer Science, vol. 15535, Springer, Cham, 2026. doi: 10.1007/978-3-032-15538-2_29.

[17] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in *Advances in Neural Information Processing Systems*, 2012. Available: `https://papers.nips.cc/paper/2012/hash/05311655a15b75fab86956663e1819cd-Abstract.html`.

[18] L. Ali, *Gestionnaire d'infrastructure distribuée*. Ph.D. dissertation, Institut National des Sciences Appliquées de Lyon, 2008.

[19] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in *Military Communications and Information Systems Conference (MilCIS)*, 2015. doi: 10.1109/MilCIS.2015.7348942.

[20] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proceedings of ICISSP*, 2018. doi: 10.5220/0006639801080116.

[21] N. Moustafa, "The BoT-IoT dataset," *Data in Brief*, vol. 24, p. 103386, 2019. doi: 10.1016/j.dib.2019.103386.