Evaluating AI/ML Techniques to Reduce False-Positives in SIEM Platforms for Resource-Constrained Environments

Student *Amit Bora* *Email:C00326758@setu.ie*

Supervisor *Mark Cummins*

Objectives

- Compare lightweight techniques: TinyML, Ensemble methods, and Adaptive thresholding.
- Evaluate models: Decision Trees, Random Forests, SVM, and Gradient Boosting.
- Maintain detection accuracy above 95% while reducing resource overhead.
- Develop a validated framework with implementation guidelines for edge deployment.

Introduction

- SIEM systems suffer from high false-positive rates (30-80%), causing alert fatigue.
- Current mechanisms are too resource-intensive for IoT and Edge devices.
- Goal: Achieve 60-85% reduction in FPR (target < 5%) with >95% detection accuracy.
- Constraint: Performance on devices with < 512MB RAM and limited CPU.

Problem Statement

- Alert Fatigue: Massive volumes of false alarms overwhelm security analysts.
- Resource Incompatibility: Standard SIEM tools exceed the memory and power limits of IoT/Edge sensors.
- Security Gap: High FPR reduces overall threat-detection effectiveness in decentralized networks.

How to deploy AI/ML-based threat detection on constrained devices while achieving low false positives, high accuracy, and operational efficiency?

Methodology

- Datasets: NSL-KDD and CICIDS2017 (simulated cybersecurity environments).
- Simulated Environment: Testing on representative boards like Raspberry Pi.
- Process: Dataset preparation → Baseline establishment → Model training → Deployment & testing → Statistical analysis.
- Metrics: False-Positive Rate (FPR), Accuracy, CPU/RAM usage, and Operating Performance.

Key Contributions

- Edge Security: Advances threat detection specifically for IoT and edge computing.
- Operational Efficiency: Directly addresses the "alert fatigue" problem for SOC analysts.
- Standardization: Provides a framework for deploying AI-enhanced components on constrained hardware.

References

- Arcot et al. (2024): TinyML for On-Device Threat Detection.
- Kapera & Niemiec (2025): Dynamic Risk Thresholds for SIEM.
- Hamidouche et al. (2024): Real-time Strategies for Constrained Devices.