Differential Privacy Trade-offs in Federated Malware Detection for Heterogeneous IoT Networks

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does the integration of differential privacy in federated learning-based malware detection models affect the trade-off between model accuracy and communication efficiency, measured by F1-score. This work investigates the possibilities enabled by federated learning concerning IoT malware detection and studies security issues inherent to this new learning paradigm. In this context, a framework that uses federated learning to detect malware affecting IoT devices is. 6 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Federated Learning for Malware Detection in IoT Devices. Research question: How does the integration of differential privacy in federated learning-based malware detection models affect the trade-off between model accuracy and communication efficiency, measured by F1-score and bandwidth usage across heterogeneous IoT device networks?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.8/10.

3 Results

10 papers retrieved. 6 claims extracted; 0 independently verified. Quality review score: 2.8/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
Federated Learning enables data privacy by design as data is not shared with any external identity.	×	0.08
Previous works dealing with FL for intrusion detection lack the use of realistic datasets in the FL context.	×	0.06
The proposed framework covers both anomaly detection and classification approaches using multi-model architectures.	×	0.08
The benchmark table on page 12 shows 95% accuracy for benign device detection in centralized models.	×	0.03
Data preprocessing is performed locally on each client device before model training.	×	0.04
Model parameters are exchanged and aggregated to create a global model instead of sharing raw data.	×	0.06

References

- http://arxiv.org/abs/2104.09994v3
- http://arxiv.org/abs/2211.12874v1
- http://arxiv.org/abs/2601.09933v1