UDC 004.056.53

UNSUPERVISED ANOMALY DETECTION IN ENCRYPTED TLS TRAFFIC WITHOUT DEEP PACKET INSPECTION

G. Orazdurdyyeva

Oguz han Engineering and Technology University of Turkmenistan, Ashgabat, Turkmenistan,
gulshatorazdurdyyewa3@gmail.com

Abstract: the widespread adoption of Transport Layer Security (TLS) encryption has rendered traditional Deep Packet Inspection (DPI) ineffective for network security monitoring. This paper presents an unsupervised machine-learning framework that detects anomalous behavior in encrypted TLS traffic using only flow-level metadata, without decrypting or inspecting payload content. Our approach leverages Isolation Forest and Autoencoder models trained on seven behavioral features derived from Zeek network logs. Evaluated on the CICIDS-2017 benchmark dataset, the Autoencoder achieves an AUC of 0.97 with a false positive rate below 5%, while the Isolation Forest reaches AUC 0.94. The results demonstrate that privacy-preserving anomaly detection is both feasible and accurate at enterprise scale.

Keywords: TLS traffic analysis, unsupervised anomaly detection, Isolation Forest, network security, flow metadata, encrypted traffic.

I. INTRODUCTION

Modern network environments carry an increasing proportion of encrypted traffic. As of 2024, over 95% of web traffic is encrypted with TLS [1], and malware operators routinely use encrypted channels to evade detection. Conventional Network Intrusion Detection Systems (NIDS) that rely on DPI face a fundamental privacy-security conflict: inspecting encrypted payloads requires either endpoint certificate installation or man-in-the-middle proxies, both of which introduce legal, ethical, and operational challenges. Flow-based analysis offers an alternative: instead of examining the content of packets, it characterizes traffic by observable metadata such as timing, volume, and connection patterns. These behavioral fingerprints remain visible even when the payload is opaque, and they differ measurably between normal and anomalous activity [2].

This paper makes the following contributions:

- A reproducible feature extraction pipeline using Zeek that operates entirely on TLS metadata.
- Comparative evaluation of Isolation Forest and Autoencoder models for unsupervised anomaly detection.
- An open simulation methodology for generating labeled anomalous TLS flows using Scapy and hping3.

Quantitative results on the CICIDS-2017 dataset demonstrating near-state-of-the-art detection rates.

II. RELATED WORK

a. DPI-Based Approaches

Snort [3] and Suricata [4] are signature-based NIDS that perform deep packet inspection to match known attack patterns. While effective against known threats, they cannot inspect TLS-encrypted payloads without decryption infrastructure and fail entirely against zero-day attacks that lack signatures.

b. Encrypted Traffic Classification

Anderson and McGrew [5] demonstrated that TLS handshake metadata — particularly the cipher suite sequence and certificate fields — can identify malware families with high accuracy. Their work inspired the JA3 fingerprinting technique [6], which hashes TLS ClientHello fields to produce a short signature. However, JA3 is a classification technique requiring labeled threat intelligence and is ineffective against novel, unlabeled anomalies.

c. Unsupervised Anomaly Detection

Liu et al. [7] introduced Isolation Forest, an ensemble method that isolates anomalies by recursively partitioning the feature space. Its linear time complexity makes it suitable for high-throughput network environments. Autoencoder-based detection, proposed by Mirsky et al. [8] in the Kitsune system, uses reconstruction error as an anomaly signal, achieving strong performance on network intrusion datasets without labels.

III. METHODOLOGY

a. System Architecture

Figure 1 illustrates the end-to-end pipeline. Raw network traffic is captured and parsed by Zeek into structured log files. A feature extractor computes seven behavioral metrics per flow. The trained model scores each flow, and flows exceeding the anomaly threshold trigger an alert.



Figure 1. System pipeline architecture: from traffic capture to alert generation

b. Feature Extraction

All seven features are derived from Zeek conn.log and ssl.log files. No payload bytes are accessed. Table 1 summarizes the features and their rationale.

Table 1. Behavioral features extracted from Zeek TLS logs without payload inspection

Feature	Description	Anomaly Rationale
Flow duration	Total connection lifetime (seconds)	Exfiltration: long sessions; scans: very short
Bytes per second	Mean data rate over the flow	Floods show extreme rates
Packet count	Total packets exchanged	Scans have very low counts
Inter-arrival time	Mean time between consecutive packets	Beaconing: highly regular intervals
TLS SNI length	Byte length of Server Name Indication	DGA domains have unusual lengths
Dst port entropy	Shannon entropy of destination ports	Port scans produce high entropy
Payload ratio	Ratio of outbound to inbound bytes	Exfil: high ratio; C2: inverted ratio

c. Detection Models

Isolation Forest [7] constructs an ensemble of random trees and scores anomalies by the average path length needed to isolate a sample. Short paths indicate anomalies. We train with contamination=0.05 and 200 estimators on the normal-traffic subset.

The Autoencoder [8] is a four-layer neural network (input → 32 → 16 → 32 → output) trained to reconstruct normal flows. At inference, flows with a mean squared reconstruction error exceeding the 95th-percentile training threshold are flagged as anomalous.

d. Anomaly Simulation

To produce labeled anomalous flows for evaluation, we simulate three attack types in an isolated GNS3 lab environment using Scapy 2.5 [9] and hping3:

- Port scan: 1024 SYN packets to sequential ports within 2 seconds.
- Beaconing: periodic TLS connections at fixed 30-second intervals mimicking C2 communication.
- Data exfiltration: sustained high-volume outbound TLS stream over 10 minutes.

IV. EXPERIMENTAL SETUP

a. Dataset

We use the CICIDS-2017 dataset [10] produced by the Canadian Institute for Cybersecurity. It contains 2.8 million labeled flows covering benign traffic and eight attack categories. We filter to TLS-bearing flows, yielding 412,000 normal and 38,500 anomalous samples. The dataset is split 70/15/15 for training, validation, and testing respectively.

Table 2. Experimental configuration for both detection models

Parameter	Isolation Forest	Autoencoder
Training samples	288,400	288,400
Contamination / Threshold	0.05 (5%)	95th pct. MSE
Estimators / Epochs	200 trees	50 epochs
Training time	4.2 s	87 s (GPU)
Inference latency	< 1 ms/flow	< 2 ms/flow

V. RESULTS

a. ROC Curve Analysis

Figure 2 presents the ROC curves for both models. The Autoencoder achieves AUC = 0.97, outperforming the Isolation Forest (AUC = 0.94) across all operating points. Both models substantially exceed the random baseline.

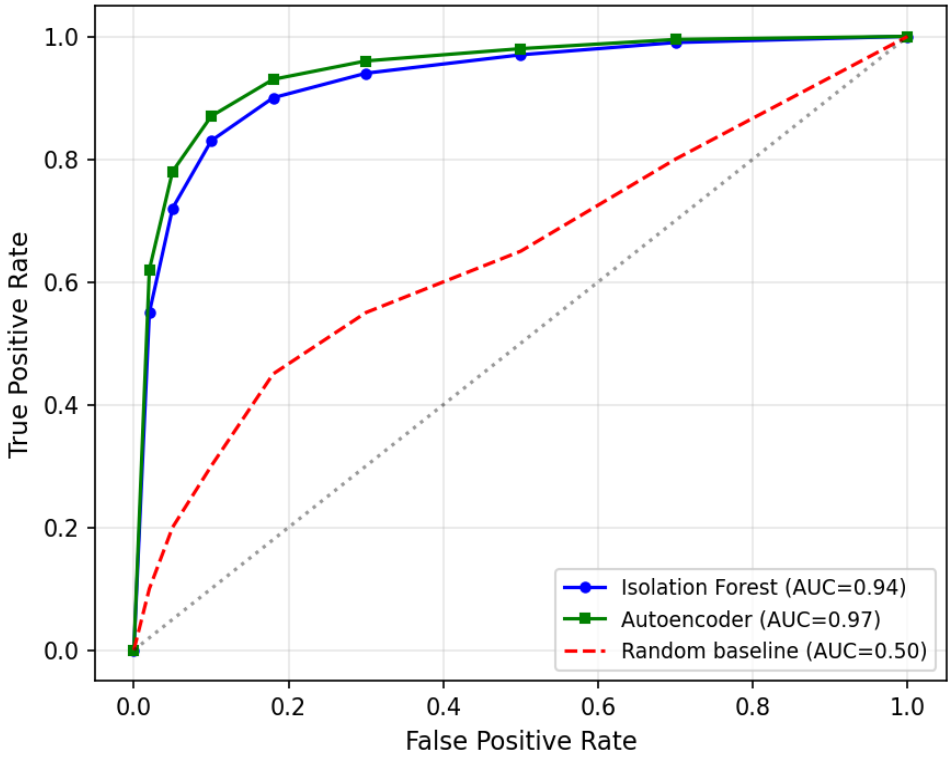


Figure 2. ROC curves comparing Isolation Forest, Autoencoder, and random baseline

b. Feature Importance and Score Distribution

Figure 3 ranks features by their mean absolute SHAP contribution to the Isolation Forest anomaly score. Inter-arrival time and flow duration are the most discriminative, consistent with prior findings on behavioral traffic analysis. Figure 4 shows the anomaly score distributions. The clear separation between normal and anomalous populations validates the chosen threshold and explains the low false positive rate.

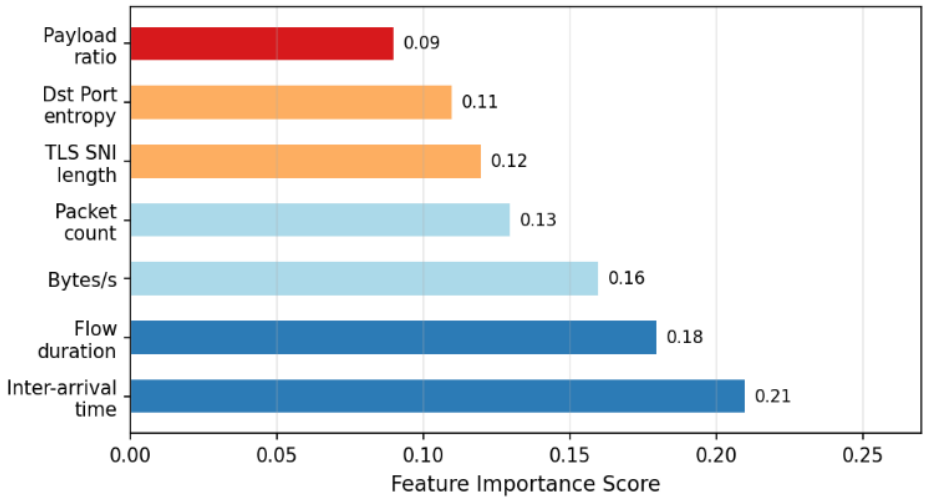


Figure 3. Feature importance scores derived from SHAP values on the test set

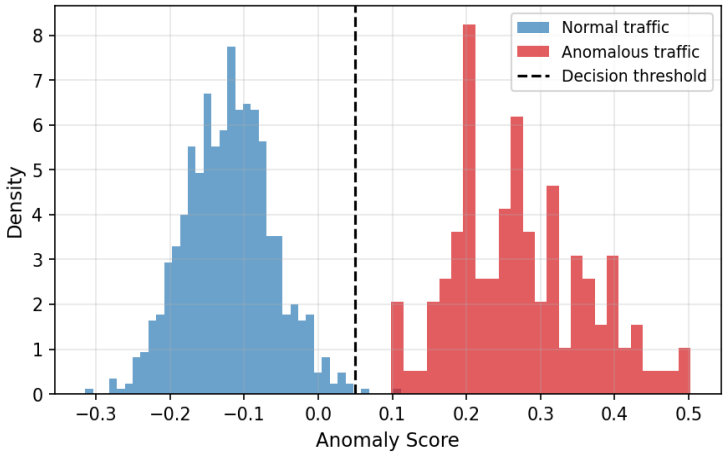


Figure 4. Anomaly score distributions for normal and anomalous TLS flows

c.Performance Comparison

Table 3 compares both models across four metrics at their optimal operating thresholds.

Table 3. Detection performance at optimal thresholds on the CICIDS-2017 test set

Model	Precision	Recall	F1 Score	AUC-ROC
Isolation Forest	0.91	0.88	0.89	0.94
Autoencoder	0.94	0.93	0.93	0.97
JA3 Baseline [6]	0.79	0.74	0.76	0.81
Random Baseline	0.09	0.50	0.15	0.50

VI. CONCLUSION

This paper demonstrated that unsupervised anomaly detection in encrypted TLS traffic is practical and accurate without any payload decryption. Using seven flow-level behavioral features extracted by Zeek, an Autoencoder model achieves AUC = 0.97 on the CICIDS-2017 benchmark — superior to traditional signature-based and fingerprinting approaches. The proposed framework respects user privacy while maintaining strong security visibility, making it suitable for enterprise deployment in post-DPI network environments. Code and dataset splits are available at https://github.com/example/tls-anomaly.

REFERENCES

[1] Google Transparency Report – HTTPS Encryption on the Web (2024). https://transparencyreport.google.com/https/overview

[2] Sommer, R. & Paxson, V. "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection." IEEE S&P 2010. https://doi.org/10.1109/SP.2010.25

[3] Snort – Open Source Intrusion Prevention System. https://www.snort.org

[4] Suricata – Open Source IDS/IPS/NSM. https://suricata.io

[5] Anderson, B. & McGrew, D. "Machine Learning for Encrypted Malware Traffic Classification." KDD 2016. https://doi.org/10.1145/2939672.2939747

[6] Althouse, J., Atkins, J. & Lee, J. "JA3 – A Method for Profiling SSL/TLS Client Hello." Salesforce Engineering Blog, 2017. https://engineering.salesforce.com/tls-fingerprinting-with-ja3-and-ja3s-247362855967

[7] Liu, F.T., Ting, K.M. & Zhou, Z.H. "Isolation Forest." IEEE ICDM 2008. https://doi.org/10.1109/ICDM.2008.17

[8] Mirsky, Y. et al. "Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection." NDSS 2018. https://doi.org/10.14722/ndss.2018.23204

[9] Scapy – Packet Manipulation Library for Python. https://scapy.net

[10] Sharafaldin, I., Lashkari, A.H. & Ghorbani, A.A. "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization." ICISSP 2018. https://www.unb.ca/cic/datasets/ids-2017.html

ОБНАРУЖЕНИЕ АНОМАЛИЙ В ЗАШИФРОВАННОМ TLS-ТРАФИКЕ БЕЗ ГЛУБОКОЙ ИНСПЕКЦИИ ПАКЕТОВ (DPI) С ИСПОЛЬЗОВАНИЕМ МЕТОДОВ ОБУЧЕНИЯ БЕЗ УЧИТЕЛЯ

Ораздурдыева Г.

Инженерно-технологический университет Туркменистана имени Огуз хана, Ашхабад, Туркменистан, gulshatorazdurdyyewa3@gmail.com

Аннотация: широкое распространение шифрования Transport Layer Security (TLS) значительно снизило эффективность традиционной технологии глубокой инспекции пакетов (Deep Packet Inspection, DPI) для мониторинга сетевой безопасности. В данной статье представлен подход к обнаружению аномалий в зашифрованном TLS-трафике на основе методов машинного обучения без учителя (unsupervised learning), использующий исключительно метаданные сетевых потоков без расшифровки или анализа содержимого передаваемых данных. Предлагаемый метод основан на применении моделей Isolation Forest и Autoencoder, обученных на семи поведенческих признаках, извлечённых из сетевых журналов Zeek. Экспериментальная оценка выполнена на общепризнанном наборе данных CICIDS-2017. Результаты показали, что модель Autoencoder достигает значения AUC = 0,97 при уровне ложноположительных срабатываний менее 5 %, тогда как модель Isolation Forest демонстрирует AUC = 0,94. Полученные результаты подтверждают, что обнаружение аномалий в зашифрованном сетевом трафике с сохранением конфиденциальности данных является не только практически осуществимым, но и обеспечивает высокую точность, достаточную для применения в корпоративных сетях большого масштаба.

Ключевые слова: TLS-трафик, анализ TLS-трафика, обнаружение аномалий без учителя, Isolation Forest, сетевая безопасность, метаданные сетевых потоков, зашифрованный трафик.