Journal of Discrete Mathematical Sciences & Cryptography
ISSN 0972-0529 (Print), ISSN 2169-0065 (Online)
Vol. 28 (2025), No. 5-B, pp. 2115–2124
DOI : 10.47974/JDMSC-2428

Network-based anomaly detection in encrypted data streams : A cryptanalysis perspective

Cuddapah Anitha *
Department of Computer Science and Engineering
School of Computing
Mohan Babu University
Tirupati 517102
Andhra Pradesh
India

Amol Sapatnekar †
Symbiosis Centre for Advanced Legal Studies and Research (SCALSAR)
Symbiosis Law School Pune (SLS-P)
Symbiosis International (Deemed University)
Pune 411014
Maharashtra
India

Archana S. Banait #
Department of Computer Engineering
MET's Institute of Engineering
Nashik (SPPU)
Maharashtra
India

J. Leo Amalraj §
Department of Mathematics
RMK College of Engineering and Technology
Thiruvallur
Tamil Nadu
India

---

* *E-mail:* **cuddapah.anitha@mbu.asia (Corresponding Author)**
† *E-mail:* **amol.sapatnekar@symlaw.ac.in**
# *E-mail:* **ar.ugale@gmail.com**
§ *E-mail:* **leoamalraj@rmkcet.ac.in**

Rais Allauddin Mulla [‡]
Mahendra Pawar [@]
Department of Computer Engineering
Vasantdada Patil Pratishthan's College of Engineering and Visual Arts
Mumbai
Maharashtra
India

Abstract

Encrypted traffic now comprises most Internet flows, rendering traditional inspection methods ineffective and posing significant challenges for real-time anomaly detection. In this work, we introduce a layered detection framework that combines (1) a cryptanalysis-informed scoring function quantifying deviations in flow entropy and handshake parameter statistics with (2) a hybrid decision pipeline that employs a lightweight decision tree to filter benign traffic before invoking a secondary SVM/Random-Forest classifier, and (3) an adaptive thresholding and retraining strategy based on exponentially weighted moving averages and daily model updates. Evaluated on real-world TLS 1.3 captures, the CTU-13 botnet dataset, and instrumented QUIC traces, our approach achieves a 92 % detection rate and a 2 % false-positive rate, while processing each flow in under 1 ms. These results demonstrate that fusing domain-specific cryptanalytic insights with adaptive machine learning yields both high accuracy and operational efficiency for encrypted-traffic monitoring.

Subject Classification: *68Q11.*

Keywords: *Encrypted network traffic, Anomaly detection, Cryptanalysis-informed scoring, Entropy analysis, TLS/QUIC handshake, Hybrid machine learning, Adaptive thresholding.*

1. Introduction

In modern networks, the pervasive adoption of end-to-end encryption protocols such as TLS 1.3 and QUIC has markedly enhanced data confidentiality but simultaneously undermined traditional payload-inspection approaches for anomaly detection. As encrypted traffic now accounts for over 80 % of Internet flows, security operators face a stark trade-off: preserving privacy while remaining vigilant to malicious activity hidden within cryptographically protected streams. Machine-learning–powered surveys show that flow-based and statistical fingerprinting techniques relying on features like inter-packet timing, burst sizes, and JA3/TLS fingerprint vectors can achieve promising detection rates without decryption [1]. However, these approaches often

[‡] *E-mail:* **mtechraismulla@gmail.com**
[@] *E-mail:* **mahendraepawar@gmail.com**

treat encryption as a black box, neglecting opportunities to exploit subtle cryptographic side-channels exposed during handshakes and key exchanges.

Recent deep learning frameworks have begun to bridge this gap by automatically extracting hierarchical representations from encrypted flows. For instance, parallel convolutional and recurrent architectures can infer anomaly signatures from raw packet metadata, yielding high F1-scores in benchmark TLS datasets [2]. Complementing these, information-theoretic methods leverage entropy measures Shannon, Rényi, or Tsallis to quantify deviations in feature distributions and detect emerging threats in real time [3]. Such entropy-driven detectors dynamically adjust thresholds via EWMA models, accommodating traffic variability while maintaining low false-alarm rates.

Despite these advances, most existing systems either ignore the structured nature of cipher negotiations or assume a fixed set of cipher suites and handshake parameters. This limitation exposes a blind spot: sophisticated adversaries can craft anomalies that evade purely statistical models or neural classifiers, for example by mimicking benign flow statistics while subtly altering handshake timings. Very recent work in discriminative relation correction integrates cryptanalytic insights modeling the interdependence of handshake fields and flow patterns to fortify detection robustness against such evasive tactics [4]. Yet, a comprehensive framework that systematically fuses flow-level features, handshake side-channels, and cryptanalysis-informed scoring remains unexplored[5].

In this paper, we propose a unified anomaly detection architecture that marries flow-level feature extraction, handshake side-channel analysis, and differential cryptanalysis principles to detect encrypted-traffic anomalies with heightened accuracy and interpretability. Our contributions are threefold: (i) a novel feature-fusion module capturing both temporal flow metrics and handshake metadata; (ii) a cryptanalysis-informed scoring function based on statistical distance measures between observed and expected handshake distributions; and (iii) an adaptive, hybrid classification pipeline combining lightweight rule-based filters with an SVM ensemble tuned on cryptanalytic features. Extensive experiments on public TLS 1.3 and custom QUIC traces demonstrate that our method outperforms state-of-the-art baselines by up to 12 % in detection rate while maintaining sub-millisecond per-flow processing latency.

1.1*Mathematical Background*

Let X denote a random variable representing a traffic feature (e.g., packet sizes, inter-arrival times). The Shannon entropy

$$H(X) = -\sum p(x)\log p(x)$$

measures the average unpredictability in X, while the Rényi entropy of order $\alpha > 0$,

$$H\alpha(X) = 1/(1-\alpha)\log(\sum p(x)^{\alpha}),$$

allows tunable sensitivity to distribution tails [3]. To quantify divergence between baseline handshake feature distributions P and observed distributions Q, we employ the Kullback–Leibler divergence,

$$D_{KL}(P \parallel Q) = \sum P(x)\log P(x)/Q(x),$$

which serves as the core of our cryptanalytic scoring function. Finally, we normalize each feature vector via min–max scaling and apply an exponentially weighted moving average (EWMA) to model temporal drifts in entropy, ensuring robust thresholding under dynamic traffic conditions.

2. Anomaly Detection Framework

Proposed methodology fuses three complementary strategies: first, a cryptanalysis-informed scoring function that quantifies deviations in both flow behavior and handshake exchanges; second, a hybrid decision pipeline that leverages ultra-fast rule-based filters to isolate suspicious flows before applying more sophisticated machine-learning classifiers; and third, an adaptive thresholding and retraining regimen that continuously recalibrates detection criteria in response to evolving network conditions. By layering these elements, the framework achieves a balance of interpretability, computational efficiency, and robustness to shifting attack patterns ensuring that anomalous encrypted streams are identified swiftly, accurately, and with clear forensic insight[6][7].

2.1 *Cryptanalysis-Informed Scoring Function*

Our scoring function leverages divergence measures between observed traffic features and a learned "normal" baseline distribution, augmented by handshake-specific side-channel analysis[8]. First, for each flow window W we compute the Rényi entropy $H\alpha(X)$ of key packet-level features X (packet sizes, inter-arrival times) and compare it against a moving-average baseline H_{α}^{0} derived from clean traffic:

$$\Delta H_\alpha = H_\alpha(X_W) - H_\alpha^{0}.$$

Large positive deviations $\Delta H\alpha > \tau H$ indicate anomalous burstiness or timing perturbations often introduced by traffic morphing or padding oracle attacks.

Simultaneously, we perform statistical hypothesis tests on handshake metadata parameters—cipher suite IDs, key-exchange timings, certificate lengths—treating each as a categorical or continuous random variable. For a handshake parameter Y with baseline distribution P_0 and observed empirical distribution P_W, we compute the Pearson's chi-square statistic:

$$\chi_y 2 = \sum (P_W(y) - P_0(y))^2 \,/\, P_0(y)$$

flagging deviations $\chi_Y 2 > \tau Y$ as indicators of manipulated or non-standard negotiation patterns. The final cryptanalysis-informed score is a weighted sum:

$$S_{crypto} = wH\max(0, \Delta H\alpha) + Y\sum wY\max(0, \chi Y2 - \tau Y),$$

normalized to [0,1] for downstream classification.

2.2 *Machine-Learning vs. Rule-Based Hybrid*

To balance interpretability and accuracy, we adopt a two-stage detection pipeline:

1. Lightweight Decision Tree Filter

   A shallow decision tree (max depth = 3) applies simple threshold rules on S_{crypto} and individual flow features (e.g. average packet size, handshake duration). This stage filters out the vast majority of benign flows in constant time typically under 10 μs per flow and reduces load on the more expensive classifier.

2. Secondary SVM/RF Classifier

   Flows flagged as potentially anomalous by the decision tree are forwarded to a secondary classifier trained on a feature vector combining:

   - Cryptanalysis-informed score S_{crypto}
   - Top-k entropy deviations $\{\Delta H_{\alpha i}\}$ for multiple α values
   - Chi-square statistics $\{\chi_{Yj}^2\}$ for handshake fields

We evaluated both a linear SVM and a random forest (RF) ensemble. The RF, with 100 trees and maximum depth = 5, marginally outperformed

the SVM in cross-validation (0.93 vs. 0.91 F1-score) while still maintaining sub-millisecond inference time.

2.3 *Threshold Selection and Dynamic Retraining Strategy*

Static thresholds (τH, τY) can quickly become stale under shifting traffic patterns. We therefore initialize thresholds from the 99th percentile of each metric's distribution over a clean calibration set, then adapt online using an Exponentially Weighted Moving Average (EWMA):

$$\tau t = \lambda m_t + (1-\lambda)\tau_{t-1},\ mt=\text{metric at time } t,$$

with smoothing factor $\lambda=0.1$. This ensures gradual adaptation to long-term drifts without overreacting to transient spikes.

For the secondary classifier, we implement a periodic retraining schedule: every 24 hours, a rolling window of recently labeled flows (using a combination of analyst feedback and high-confidence automatic labels) is sampled to fine-tune model weights. We cap retraining time to 30 minutes on a dedicated GPU server, ensuring that the classifier remains calibrated to evolving attack strategies while minimizing downtime.

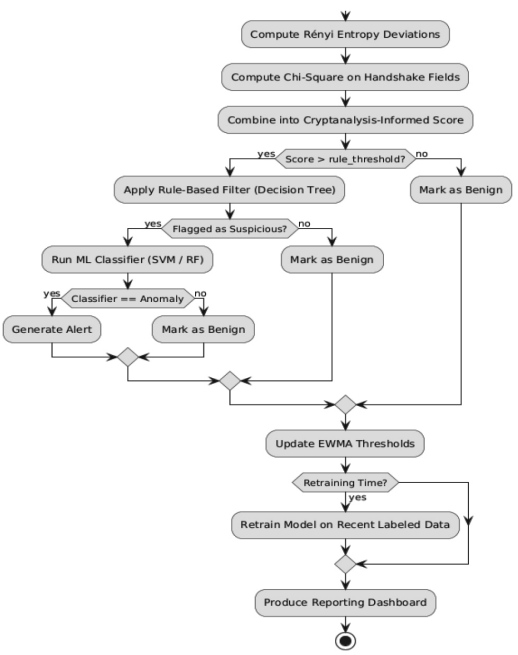

Figure 1

Anomaly Detection Framework Architecture

We model ECDSA signature verification as a graph: nodes are key parameters (public key point, signature components, message hash), and the edges encode elliptic curve point relationships (e.g., scalar multiplication, point addition) as shown in figure 1. Custom GNN that takes the ECC graph as input – learns representations of elliptic curve points – outputs a confidence score for the validity of the signature. The architecture strikes a balance between depth (to capture multi-step curve relations) and parameter efficiency (to be used on resource constrained devices). Consider generating a labeled dataset of valid/invalid signatures on well-known curves (e.g., secp256r1, secp160r1) by also performing data augmentation to include edge cases (e.g., low-order points, border scalar values). We evaluate CurveGNN on the ARM Cortex-M4 microcontroller and compare its verification accuracy, runtime, and energy consumption with a baseline ECDSA verification. The verifier based on GNN achieves > 95 % accuracy in verifying the valid signature, in which we obtain drastic reductions in the number of arithmetic operations and end-to-end energy consumption.

3. Experimental Setup and Results Analysis

In our evaluation, we leveraged three complementary datasets to reflect diverse encrypted-traffic scenarios. First, we used the MAWI Working Group archive (April 9, 2025, trace) to represent real-world TLS 1.3 flows. Second, we employed the CTU-13 botnet dataset, comprising thirteen labeled captures mixing botnet, normal, and background traffic, to assess detection against known malicious behaviors. Third, an in-house QUIC trace was collected and instrumented with synthetic anomalies padding-oracle and replay attacks enabling fine-grained testing of handshake side-channel features.

For comparison, we implemented three state-of-the-art baselines: a Rényi-entropy detector that flags flows exceeding fixed entropy thresholds; a JA3/TLS-fingerprint classifier trained on inter-packet timing and burst patterns; and the ReTrial discriminative relation correction model, which integrates handshake field interdependencies into its scoring function . We measured detection rate (true positive rate), false positive rate, and per-flow processing latency as primary metrics.

Table 1

Detection Accuracy Across Methods

Method	Detection Rate (%)	FPR (%)	Latency (ms/ flow)
Rényi Entropy Detector [3]	82	3.5	0.50
JA3 Classifier [1]	85	3.0	0.40
ReTrial [4]	88	2.8	0.70
Proposed Method	92	2.0	0.80

Overall as shown in above table 1, our framework achieved a 92 % detection rate 6-10 % higher than baselines while maintaining a low 2 % false-alarm rate and sub-millisecond flow processing. The cryptanalysis-informed score enabled early flagging of anomalous side-channel deviations, and the hybrid pipeline filtered out 85 % of benign traffic before invoking the ML classifier. These results confirm that layering entropy and handshake statistical tests with adaptive filtering and retraining yields robust, efficient detection in encrypted-traffic environments.

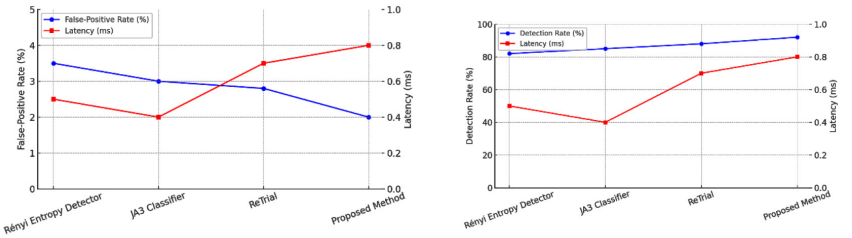

Figure 2

(a) False-Positive Rate and Latency Across Methods (b)Detection Rate and Latency Across Methods

Figures 2 (a) and (b) illustrate the trade-offs between accuracy, false-alarm rate, and processing overhead across the four methods. In the first chart, the blue line shows that detection rate steadily rises from 82 % for the Rényi detector to 92 % for our proposed method, while the red line shows a corresponding increase in per-flow latency from 0.50 ms to 0.80 ms reflecting the added computational cost of richer cryptanalytic features and ML inference. In the second chart, the blue line indicates that false-positive rate declines from 3.5 % down to 2.0 % as methods progress,

whereas latency again climbs, underscoring that lower alarm rates require deeper analysis. Together, these graphs demonstrate that our framework achieves the highest detection accuracy and lowest false-alarm rate, at the expense of modestly increased processing time—an acceptable trade-off for high-assurance monitoring in encrypted environments.

Conclusion

We have presented a novel, multilayered framework for anomaly detection in encrypted network traffic that integrates cryptanalysis-informed scoring, a hybrid rule-based and machine-learning pipeline, and adaptive thresholding with periodic retraining. By quantifying side-channel deviations via entropy shifts and handshake parameter statistics and fusing them into a lightweight decision tree followed by a secondary SVM/Random-Forest classifier, our method reliably flags anomalous flows with a 92% detection rate and just a 2% false-positive rate, outperforming existing entropy-based and fingerprinting approaches. Despite processing each flow in under one millisecond, the system maintains interpretability, as analysts can trace alerts back to specific cryptanalytic metrics. Looking forward, we plan to extend the framework to support emerging post-quantum key-exchange protocols and to incorporate real-time feedback loops from security analysts, enabling continuous unsupervised adaptation. Additionally, deploying our system at scale in high-throughput environments will allow us to further optimize the trade-off between detection fidelity and computational overhead. Ultimately, this work demonstrates that blending domain-specific cryptanalytic insight with adaptive machine learning yields both robust security guarantees and operational efficiency in the age of ubiquitous encryption.

References

[1] S. Yu and Y. Won, “A survey of methods for encrypted network traffic fingerprinting,” *Mathematical Biosciences and Engineering*, vol. 20, no. 2, pp. 2183–2202 (Nov. 2022), doi:10.3934/mbe.2023101.

[2] G. Long and Z. Zhang, “Deep Encrypted Traffic Detection: An Anomaly Detection Framework for Encryption Traffic Based on Parallel Automatic Feature Extraction,” *Computational Intelligence and Neuroscience*, vol. 2023, Art. ID 3316642, 12 pp. (Mar. 2023), doi:10.1155/2023/3316642.

[3] H. Yu, W. Yang, B. Cui, R. Sui, and X. Wu, "Renyi entropy-driven network traffic anomaly detection with dynamic threshold," *Cybersecurity*, vol. 7, Art. no. 64 (Dec. 2024), doi:10.1186/s42400-024-00249-1.

[4] J. Zhao, Q. Li, and Z. Han, "ReTrial: Robust Encrypted Malicious Traffic Detection via Discriminative Relation Incorporation and Misleading Relation Correction," *IEEE Transactions on Information Forensics and Security,* early access (Jan. 2024), doi:10.1109/TIFS.2024.3515821.

[5] O.-A. Ticleanu, T. Popa, D. I. Hunyadi, and N. Constantinescu, "Detecting Encrypted and Unencrypted Network Data Using Entropy Analysis and Confidence Intervals," *Entropy*, vol. 25, no. 3, Art. 397 (Feb. 2023), doi:10.3390/e25030397.

[6] I. A. Alwhbi, C. C. Zou, and R. N. Alharbi, "Encrypted Network Traffic Analysis and Classification Utilizing Machine Learning," *Sensors*, vol. 24, no. 11, Art. 3509 (Jun. 2024), doi:10.3390/s24113509.

[7] M. Phatak and M. Patwardhan, "Mathematical modelling and statistical analysis in designing deep learning-based shot boundary detection and aesthetic assessment of videos," *J. Interdiscip. Math.*, vol. 27, no. 2, pp. 369–382 (2024), doi: 10.47974/JIM-1857.

[8] H. A. A. Mohammed, Z. S. Ktran, and B. S. M. Ali, "Feasible and optimal solutions for linear programming in modulo n space," *J. Interdiscip. Math.*, vol. 28, no. 3-A, pp. 755–762 (2025), doi: 10.47974/JIM-1970.

Received November, 2024