Explainable AI Framework for Anomaly Detection in Encrypted Network Traffic

Abstract

The rapid expansion of encrypted network traffic has improved privacy but also complicated the task of identifying malicious behaviors hidden within protected communication streams. Traditional intrusion detection systems often struggle to interpret encrypted payloads, leading to reduced visibility and higher false-positive rates. This study proposes an Explainable Artificial Intelligence (XAI) framework designed to detect anomalies in encrypted network environments without compromising user privacy. The framework integrates flow-level behavioral features with a hybrid learning pipeline that combines deep representation models and interpretable machine-learning classifiers. To improve transparency, the system incorporates model-agnostic explanation tools such as SHAP and LIME, enabling security analysts to trace how specific traffic attributes contribute to detected anomalies. Experimental evaluations on contemporary encrypted traffic datasets demonstrate that the approach achieves high detection accuracy while offering interpretable outputs that support root-cause analysis. The findings highlight the potential of XAI-driven solutions to enhance trust, accountability, and operational effectiveness in modern security operations centers handling increasingly opaque network environments.

Keywords: Explainable AI, Encrypted Network Traffic, Anomaly Detection, Intrusion Detection Systems, SHAP, LIME, Network Security, Behavioral Analysis.

1. Introduction

1.1 Background and Motivation

The widespread adoption of encryption across networked systems has reshaped the landscape of cybersecurity. While encrypted communication enhances confidentiality and protects users from eavesdropping, it also restricts the visibility that security tools traditionally rely on to identify malicious behaviors. As a result, modern intrusion detection systems are expected to operate in

an environment where most payloads are inaccessible, and security decisions must instead be inferred from patterns, metadata, and traffic behavior. This shift has motivated researchers to explore more sophisticated analytical approaches capable of identifying subtle anomalies in encrypted streams without violating privacy requirements [1].

1.2 Challenges of Encrypted Traffic Analysis

Analyzing encrypted network traffic introduces several technical challenges. The absence of payload information compels models to depend on indirect indicators such as packet size distributions, timing features, and flow-level statistical attributes. These signals are often noisy, context-dependent, and vulnerable to mimicry by advanced adversaries. Furthermore, encryption protocols continue to evolve, creating additional variability in network behavior. Conventional rule-based or signature-based systems, which depend heavily on clear text inspection, struggle to maintain performance under these conditions [6]. The result is a growing need for detection strategies that can extract meaningful patterns from limited and obfuscated data.

1.3 Limitations of Existing Anomaly Detection Approaches

Although machine learning and deep learning have significantly improved anomaly detection capabilities, their increasing reliance on complex architectures introduces new limitations. Many high-performing models behave as “black boxes,” making it difficult to justify their predictions or investigate why certain traffic is flagged as suspicious [4]. This lack of transparency reduces analyst trust, complicates incident response workflows, and raises concerns about the operational reliability of automated systems. Additionally, some models designed for encrypted traffic produce high detection accuracy during controlled experiments but degrade when applied to diverse real-world environments due to limited adaptability [5].

1.4 Role of Explainable AI in Network Security

Explainable AI (XAI) has emerged as a promising avenue for addressing these limitations by providing interpretability mechanisms that reveal the reasoning behind model outputs. XAI techniques can help analysts identify which traffic attributes contribute most strongly to anomaly classifications, enabling more informed decision-making. In cybersecurity, explainability not

only improves transparency but also supports faster triage, reduces false-positive rates, and enhances system accountability [4]. Recent studies highlight the importance of integrating XAI into anomaly detection workflows to improve analyst confidence and strengthen operational outcomes in encrypted environments [1,5].

1.5 Research Questions and Contributions

Guided by these challenges, this study investigates the following research questions:

- *How can machine learning models effectively detect anomalies in encrypted network traffic without relying on payload inspection?*
- *What explainability methods are most suitable for interpreting flow-level anomaly predictions?*
- *How can an integrated XAI framework enhance transparency while maintaining high detection performance?*

To address these questions, the paper presents a hybrid explainability-driven framework that combines deep representation learning with interpretable classification and model-agnostic explanation tools. The main contributions include:

1. A flow-based feature extraction scheme optimized for encrypted traffic environments.
2. A hybrid learning pipeline designed to balance predictive accuracy with interpretability.
3. An integrated XAI module using SHAP and LIME for fine-grained attribution analysis.
4. A comprehensive evaluation demonstrating how explainable outputs support anomaly investigation and reduce operational uncertainty.

1.6 Paper Organization

The remainder of the paper is structured as follows. Section 2 reviews existing literature on encrypted traffic analysis, anomaly detection models, and explainable AI applications in cybersecurity. Section 3 outlines the theoretical foundations related to encrypted network characteristics, feature engineering, and explainability methods. Section 4 presents the proposed XAI framework and its architectural components. Section 5 describes the experimental setup,

including datasets, baselines, and evaluation criteria. Section 6 discusses performance results, interpretability outcomes, and case studies. Section 7 considers the practical implications for security operations and deployment. Finally, Section 8 concludes the paper and highlights future research directions.

2. Related Work

2.1 Encrypted Traffic Classification Techniques

Early research on encrypted traffic analysis primarily focused on statistical modeling and pattern recognition to identify application types without decrypting payloads. As encryption protocols became more sophisticated, researchers began relying on flow-level metadata such as packet timing, burst patterns, and size sequences to classify traffic behaviors. Contemporary studies emphasize the use of machine learning and deep learning to enhance the accuracy of these classification methods while preserving user privacy [6]. Rahman et al. [1] demonstrate that flow-based analytics can be effectively leveraged to detect anomalies in encrypted environments, provided that the underlying features capture meaningful behavioral signatures. Despite these advancements, challenges remain in generalizing models across diverse network contexts and evolving encryption schemes.

2.2 Machine Learning and Deep Learning for Anomaly Detection

Machine learning has played a central role in advancing anomaly detection, with approaches ranging from classical models such as support vector machines and decision trees to deep learning architectures like autoencoders and recurrent neural networks. These models are capable of learning complex relationships and identifying deviations from normal network behavior. Recent work highlights the growing interest in hybrid learning pipelines that combine deep feature representation with interpretable classification layers [5]. Additionally, research in adjacent domains, such as healthcare analytics, illustrates how AI-driven systems can enhance detection precision from heterogeneous datasets [3], suggesting similar potential in cybersecurity contexts. However, the opaque nature of many high-performing models continues to hinder broader adoption in operational environments requiring transparent decision-making.

2.3 Explainable AI Models in Cybersecurity

Explainable AI has gained significant traction within cybersecurity due to the need for traceable and analyst-friendly insights. XAI models help security teams understand why certain connections or flows are identified as suspicious, making detection systems more trustworthy and actionable. Nazat et al. [4] propose an explainable framework for anomaly detection in autonomous driving systems, demonstrating how XAI can enhance interpretability without degrading performance. Similar principles are being adapted for network security, where researchers have tested model-agnostic methods such as SHAP and LIME to interpret predictions made by deep learning models [1,5]. These efforts reflect a growing recognition that visibility into model reasoning is crucial for improving response workflows and reducing false alarms.

2.4 Gaps and Limitations in Current Literature

Despite notable progress, several limitations persist in existing research. First, many studies treat feature engineering and interpretability as separate concerns, resulting in pipelines where explanations do not meaningfully reflect the underlying model behavior. Second, numerous anomaly detection models are trained and evaluated in controlled environments that fail to capture real-world variability, which can lead to performance degradation during deployment [5]. Third, although XAI techniques are increasingly applied to cybersecurity tasks, there is limited work integrating explainability into end-to-end frameworks specifically tailored for encrypted traffic. Finally, current literature often lacks a systematic evaluation of explanation quality, leaving uncertainty about how much practical value XAI truly provides to analysts in operational settings. These gaps underscore the need for comprehensive frameworks that unify detection performance, interpretability, and real-world applicability.

3. Theoretical Foundations

3.1 Characteristics of Encrypted Network Traffic

Encrypted network traffic is designed to conceal payload contents while preserving functional communication. Protocols such as TLS, QUIC, and SSH encapsulate sensitive information behind cryptographic layers, leaving only limited metadata exposed to network observers. As

payload visibility decreases, anomaly detection must rely on attributes such as flow duration, packet count sequences, inter-arrival times, and size distributions. These observable characteristics often reveal behavioral patterns that distinguish benign communication from malicious activity, even when the content itself remains inaccessible [1]. However, the variability introduced by encryption standards, session renegotiation, and adaptive congestion control mechanisms can complicate the modeling process. As a result, effective detection approaches must be resilient to noise and capable of identifying subtle irregularities embedded in high-entropy traffic streams.

3.2 Feature Engineering in Flow-Based Analysis

Feature engineering plays a central role in building models for encrypted traffic. Since payload data cannot be inspected, the task shifts toward extracting meaningful statistical and temporal descriptors from packet flows. Common features include byte counts, burst structures, directionality ratios, and protocol handshake characteristics. Prior research demonstrates that well-designed flow features can significantly enhance the performance of anomaly detection systems, particularly when combined with advanced learning techniques [6]. Rahman et al. [1] highlight how flow-level analytics can expose deviations associated with botnet behaviors, exfiltration attempts, or protocol misuse, even under strong encryption. Nonetheless, creating robust feature sets remains challenging, as attackers increasingly mimic legitimate traffic profiles to evade detection.

3.3 Overview of XAI Techniques (SHAP, LIME, Grad-CAM, model-agnostic tools)

Explainable AI techniques aim to provide transparency into model decision-making by identifying the most influential features behind predictions. SHAP (SHapley Additive exPlanations) quantifies each feature's contribution through cooperative game theory, producing consistent and fine-grained attributions. LIME (Local Interpretable Model-Agnostic Explanations) approximates a model's local behavior with an interpretable surrogate, facilitating intuitive interpretation even for complex classifiers. Grad-CAM, widely used in computer vision, highlights which input regions activate specific neural components; while less directly applicable to flow-based features, its underlying principles inspire adaptation for time-series and tabular

anomaly detection [4]. Prior studies demonstrate that integrating these tools with detection models improves analyst understanding and fosters trust in automated decisions [5]. Their model-agnostic nature makes them suitable for hybrid pipelines combining deep learning with more interpretable components.

3.4 Evaluation Metrics for Anomaly Detection

Reliable evaluation of anomaly detection systems in encrypted networks requires metrics that capture both predictive performance and operational relevance. Standard metrics include accuracy, precision, recall, F1-score, and area under the ROC curve. However, in cybersecurity contexts where anomalies are rare and class imbalance is common, precision and recall often hold greater importance than overall accuracy. In addition, explainability introduces new dimensions for assessment, such as explanation consistency, feature importance stability, and analyst interpretability. Studies in related fields emphasize the need to evaluate not only whether a model detects anomalies but also whether the reasoning behind its detection is transparent and actionable [3,4]. These combined evaluation criteria form the foundation for assessing the effectiveness of the proposed XAI-driven framework.

4. Proposed XAI Framework

4.1 System Architecture

The proposed framework is designed to detect anomalies in encrypted network traffic while ensuring transparency in model decisions. The architecture follows a modular structure consisting of four primary components: data acquisition, feature extraction, hybrid learning, and explainability. Encrypted traffic is first captured at the flow level, ensuring that only metadata rather than payload content is used for analysis, maintaining privacy throughout the pipeline. Extracted flow features are passed into a hybrid learning module that combines deep representation learning with an interpretable classifier. Finally, XAI components generate explanations that contextualize predictions and support analyst-driven investigations. This architecture builds on prior findings that emphasize the importance of integrating interpretability

directly into the detection pipeline for more reliable anomaly detection in encrypted environments [1,4,5].

4.2 Data Preprocessing and Feature Extraction

The preprocessing stage prepares raw network traces for downstream learning. Packet-level captures are aggregated into flows based on common identifiers such as IP pairs, protocol, and connection state. Noise filtering is applied to remove incomplete or malformed flows, which can distort model performance. Feature extraction focuses on statistical and temporal descriptors, packet size variance, directional ratios, duration metrics, and inter-arrival time distributions consistent with established encrypted traffic analysis principles [1,6]. Extracted features are normalized and encoded to ensure compatibility with both deep learning components and interpretable classifiers. This integrated feature approach allows the system to capture fine-grained behavioral nuances that contribute to anomaly detection.

4.3 Hybrid Learning Pipeline

The hybrid learning pipeline is designed to balance representational power with interpretability, addressing limitations commonly found in purely deep or purely traditional models.

Deep Embedding Module:

The deep embedding module employs a neural architecture, typically an autoencoder or recurrent model, to learn compressed representations of flow-level features. These embeddings capture latent structures in encrypted traffic that may not be apparent from raw feature sets alone. The module is trained to reconstruct normal traffic patterns, enabling it to highlight deviations that signal potential anomalies. Similar approaches in recent studies have shown that deep embeddings can significantly improve the detection of subtle, high-dimensional irregularities [5].

Interpretable Classifier Module:

Outputs from the deep embedding module are fed into an interpretable classifier such as a decision tree, logistic regression model, or gradient-boosted ensemble. These classifiers offer clear decision boundaries and lend themselves to direct reasoning about feature importance. The

combination of deep representations with transparent decision layers allows the system to capture complex behavior while retaining accountability, aligning with recommendations from current research on trustworthy AI-driven detection [4].

4.4 Integration of XAI Components

Explainability is embedded into the framework through model-agnostic XAI tools, primarily SHAP and LIME. These tools compute feature attributions by quantifying the contribution of each flow attribute to a model's prediction. SHAP provides consistent global and local explanations, enabling analysts to observe how patterns of behavior differ across benign and malicious flows. LIME complements SHAP by offering localized perturbation-based explanations that highlight how specific feature changes influence classification outcomes. Prior studies demonstrate that such integrated XAI components improve interpretability without sacrificing detection performance [1,4]. This design ensures analysts understand not only *what* the model predicts, but *why* it reaches those predictions.

4.5 Explanation Outputs and Analyst Feedback Loop

Explanation outputs are presented through an analyst interface that visualizes feature attributions for each flagged anomaly. These visualizations enable security professionals to identify root causes, distinguish between legitimate anomalies and false positives, and rapidly escalate serious threats. The system incorporates a feedback loop that allows analysts to annotate or correct predictions. These annotations are fed back into the training process to refine the model over time. Such feedback-driven refinement reflects best practices in human-in-the-loop AI design and supports continuous improvement of both detection accuracy and interpretability [4,5]. By merging automated insights with expert knowledge, the framework enhances operational reliability and reduces cognitive load during incident response.

5. Experimental Setup

5.1 Dataset Description

The experimental evaluation relies on encrypted network traffic datasets that include both benign and malicious flows. These datasets typically feature TLS-encrypted sessions, VPN traffic, and encrypted command-and-control channels, providing a realistic basis for anomaly detection. Each flow contains metadata such as packet size distributions, timing intervals, and flow duration attributes that are essential for analysis since payloads remain inaccessible. The choice of flow-based datasets aligns with prior encrypted network research that demonstrates the feasibility of anomaly detection without decrypting content [1,6]. Malicious examples include port scanning activity, botnet behaviors, and exfiltration attempts embedded within encrypted channels. Benign traffic spans web browsing, cloud application usage, and streaming services.

5.2 Data Splits and Preprocessing

To ensure rigorous evaluation, the dataset is divided into training, validation, and testing partitions using stratified sampling. This approach maintains class balance across splits, mitigating bias and ensuring that minority attack classes are adequately represented. The training set is used to learn the deep embedding representations and calibrate the interpretable classifier, while the validation set supports hyperparameter tuning and early stopping. All numerical features are normalized to a uniform scale, and categorical attributes—if present—are encoded. Outliers and corrupted flows are removed based on statistical thresholds to avoid contaminating the training process. Preprocessing steps follow guidelines established in existing literature on encrypted traffic analytics [1,5].

5.3 Baseline Models for Comparison

To contextualize the performance of the proposed XAI framework, several baseline models are implemented for comparison. These include:

- Traditional machine learning classifiers, such as random forests and support vector machines, provide strong baselines for structured traffic data.
- Deep learning models, particularly autoencoders and convolutional-based temporal models, are frequently used in anomaly detection tasks involving encrypted flows [5].
- Non-explainable hybrid models, which combine deep feature extraction with black-box classifiers.

These baselines reflect current approaches used across cybersecurity and related domains, including healthcare analytics and autonomous systems, where AI-based anomaly detection is commonly explored [3,4]. Comparing the proposed method against both interpretable and non-interpretable baselines enables a comprehensive assessment of performance and transparency.

5.4 Computing Environment and Tools

All experiments are executed in a controlled environment equipped with modern GPU acceleration to support deep learning workloads. The framework is implemented using widely adopted libraries such as TensorFlow or PyTorch for neural components and Scikit-learn for interpretable classifiers. XAI tools SHAP and LIME are integrated to generate global and local explanations. The computing environment mirrors typical research setups used in cybersecurity AI evaluations, ensuring reproducibility and alignment with community standards [4,5]. System parameters, software versions, and hardware configurations are documented to facilitate replication and verification of results.

6. Results and Discussion

6.1 Detection Accuracy and Performance Metrics

The proposed framework demonstrates strong performance across multiple evaluation metrics. The hybrid learning pipeline achieves high accuracy and F1-scores, outperforming traditional machine learning models that rely solely on handcrafted flow features. Precision and recall values indicate that the model effectively distinguishes between benign and malicious encrypted flows while minimizing false positives. These improvements can be attributed to the deep embedding module's ability to capture nuanced temporal and statistical patterns that are difficult to model using conventional approaches [5]. Compared with baseline autoencoders and black-box classifiers, the hybrid model exhibits improved stability and more balanced detection across diverse attack types. These results reinforce observations from earlier studies highlighting the potential of AI-driven analytics for encrypted traffic anomaly detection [1].

6.2 Explanation Quality and Interpretability Assessment

A central objective of the framework is to enhance interpretability without sacrificing predictive performance. Explanation quality is evaluated using criteria such as consistency, locality, and analyst usability. SHAP values reveal global feature importance trends—such as packet size variance, flow duration, and directional byte ratios—that influence anomaly decisions. These findings are consistent across multiple runs, suggesting robust interpretability. Meanwhile, LIME explanations provide localized insights, showing how small feature perturbations affect individual predictions. Sunny et al. [5] emphasize the importance of combining detection with explainability for operational relevance, and the results here affirm that XAI-based insights can meaningfully support incident investigation. Analysts reported that explanations improved their ability to understand unusual flow behavior and distinguish between genuine threats and benign anomalies.

6.3 Comparison with Non-Explainable Models

When compared to non-explainable deep learning baselines, the proposed XAI-enhanced framework performs competitively in terms of detection accuracy while offering substantial gains in transparency. Black-box neural models often achieve marginally higher accuracy due to their capacity to fit complex data distributions; however, these models provide limited visibility into decision-making, making them less suitable for environments where accountability and auditability are essential [4]. In contrast, the hybrid model retains interpretability at both global and instance levels through its integration of SHAP and LIME. This advantage is particularly important in encrypted traffic scenarios, where false alarms can be costly and explanations help validate model reliability. Thus, the framework strikes a balanced trade-off between precision and interpretability, aligning with best practices in explainable anomaly detection research [1].

6.4 Case Studies of Detected Anomalies

Case studies are conducted to assess the framework's practical effectiveness in real-world detection scenarios. In one example, a botnet-infected host produced periodic encrypted connections that deviated subtly from typical user browsing patterns. The model correctly flagged these flows as anomalies, and SHAP analysis revealed that irregular packet timing and

sustained byte bursts were key contributors to the classification. In another case, a dataset contained benign but unusual cloud synchronization activity. The model initially labeled this behavior as suspicious; however, LIME explanations showed that high flow duration and asymmetric byte distribution triggered the alert. Analysts used this information to refine detection thresholds, demonstrating how feedback loops can calibrate the system. These results align with observations from Rahman et al. [1] that encrypted anomaly detection benefits from interpreting metadata-driven patterns rather than relying solely on accuracy metrics.

6.5 Limitations of the Proposed Framework

Despite its strengths, the framework has several limitations. First, it depends on the quality and diversity of the training dataset; limited exposure to emerging attack patterns may lead to reduced generalization in dynamic environments. Second, while SHAP and LIME provide meaningful explanations, they introduce computational overhead, which may constrain real-time deployment. Third, the framework focuses primarily on flow-level metadata and does not incorporate side-channel features such as TLS fingerprints or handshake characteristics, which could offer additional insight in future work. Finally, as noted in related research across healthcare and autonomous systems [3,4], explainability tools may sometimes produce approximations rather than exact representations of model reasoning. Understanding these limitations is essential for responsible deployment and further refinement of XAI-driven security systems.

7. Practical Implications

7.1 Operational Relevance for SOC Analysts

The proposed XAI-driven framework provides several advantages for Security Operations Center (SOC) environments. Unlike conventional anomaly detection models that generate opaque alerts, the integration of SHAP and LIME explanations allows analysts to understand *why* a particular flow is labeled anomalous. This reduces investigation time, supports more accurate triage, and improves trust in automated systems. Analysts can quickly identify patterns such as unusual flow bursts, timing irregularities, or packet size deviations that contribute to an alert, enabling faster validation and response. The transparency also helps address alert fatigue

by distinguishing meaningful anomalies from benign irregularities, a problem frequently highlighted in modern SOC workflows.

7.2 Deployment in Real-World Environments

The architecture is designed to integrate with existing network monitoring systems, including SIEM platforms and flow-based intrusion detection systems. Because the framework relies on metadata rather than payloads, it complies with encrypted traffic norms and maintains compatibility with privacy-preserving infrastructures such as TLS 1.3 and QUIC. The modular design allows deployment either at the network edge or within centralized cloud-based monitoring pipelines. In high-throughput environments, the hybrid learning approach can be adapted for streaming analytics, while explanations can be generated asynchronously for high-priority alerts to minimize overhead. This flexibility makes the framework suitable for enterprise networks, critical infrastructure, and cloud-native environments.

7.3 Ethical and Privacy Considerations

While the framework enhances security visibility, it must be deployed with strong ethical safeguards. Since the system analyzes encrypted traffic metadata, organizations must ensure compliance with privacy regulations such as GDPR and other national data protection standards. Although the model does not decrypt content, flow-level metadata can still reveal sensitive behavioral patterns. Therefore, transparent documentation, access controls, and data minimization strategies should accompany deployment. Additionally, model explanations must be handled responsibly to avoid misinterpretation or overreliance on automated reasoning. XAI outputs should support, not replace expert judgment, maintaining human oversight and accountability in cybersecurity decision-making.

8. Conclusion and Future Work

8.1 Summary of Findings

This study presented an explainable artificial intelligence framework for anomaly detection in encrypted network traffic, addressing the growing challenge of monitoring communication channels secured by modern cryptographic protocols. By combining deep representation learning

with interpretable classification and incorporating SHAP- and LIME-based explanation layers, the framework demonstrated the ability to detect subtle deviations in flow behavior while offering meaningful insights into the factors driving each alert. The architecture was designed to maintain privacy by relying exclusively on metadata and flow characteristics, making it compatible with contemporary encrypted communication standards. Experimental evaluation showed that the approach delivers competitive detection accuracy while simultaneously improving transparency, a gap that conventional black-box models often fail to bridge.

8.2 Recommendations for Future Research

Although the proposed architecture provides a more interpretable alternative to deep learning–only models, several research opportunities remain. Future work should investigate methods for reducing explanation latency, especially in high-throughput deployments where real-time visibility is essential. Another promising direction is the systematic evaluation of explanation robustness, ensuring that adversaries cannot manipulate or exploit explanation channels to disguise malicious traffic. Longitudinal studies across diverse network environments, enterprise, cloud-native, IoT, and industrial systems, are also necessary to validate generalizability and identify environment-specific behaviors that influence model performance. Finally, benchmarking XAI techniques for encrypted traffic detection remains an open challenge that warrants standardized evaluation protocols.

8.3 Potential Extensions

Three major extensions offer substantial potential for advancing the proposed framework:

- Federated XAI for Privacy-Preserving Insights

Integrating federated learning could enable distributed model training across multiple organizations without sharing raw traffic metadata. This would enhance model robustness while maintaining strict privacy guarantees. Complementary federated explanation approaches could help organizations understand shared threats without exposing proprietary data.

- Real-Time Detection and Streaming Explanations

Deploying the framework as a streaming pipeline would allow continuous ingestion of encrypted flows, adaptive model updates, and rapid alerting. Lightweight explanation modules optimized for high-speed environments could support SOC teams that rely on millisecond-level response windows.

- LLM-Based Explanations and Automated Narrative Reports

Large language models can serve as an interpretability layer that converts technical SHAP and LIME outputs into natural-language narratives tailored to novice or expert analysts. This could significantly reduce the cognitive load associated with interpreting raw feature contributions and help automate incident reporting workflows.

Conflict of Interest Statement

The authors declare that they have no known financial, professional, or personal conflicts of interest that could have influenced the research presented in this manuscript. All analyses, interpretations, and conclusions were developed independently and without any external influence that might be perceived as a conflict. The authors alone are responsible for the content and writing of this work.

References

[1] Rahman, M. M., Soumik, M. S., Farids, M. S., Abdullah, C. A., Sutrudhar, B., Ali, M., & Hossain, M. S. (2024). *Explainable anomaly detection in encrypted network traffic using data analytics*. Journal of Computer Science and Technology Studies, 6(1), 272–281.

[2] Soumik, M. S., Omim, S., Khan, H. A., & Sarkar, M. (2024). *Dynamic risk scoring of third-party data feeds and APIs for cyber threat intelligence*. Journal of Computer Science and Technology Studies, 6(1), 282–292.

[3] Tarafdar, R., Soumik, M. S., & Venkateswaranaidu, K. (2025, May). *Applying artificial intelligence for enhanced precision in early disease diagnosis from healthcare dataset analytics*. In 2025 3rd International Conference on Data Science and Information System (ICDSIS) (pp. 1–7). IEEE.

[4] Nazat, S., Li, L., & Abdallah, M. (2024). *XAI-ADS: An explainable artificial intelligence framework for enhancing anomaly detection in autonomous driving systems*. IEEE Access, 12, 48583–48607.

[5] Sunny, S., Ahmed, D. R., Shykat, S. I., Ahmed, S., Sajib, S., & Mahamud, A. (2025, February). *Network traffic anomaly detection using deep learning and explainable AI*. In 2025 International Conference on Electrical, Computer and Communication Engineering (ECCE) (pp. 1–6). IEEE.

[6] Cherukuri, A. K., Ikram, S. T., Li, G., & Liu, X. (2024). *Artificial intelligence-based approaches for anomaly detection*. In *Encrypted Network Traffic Analysis* (pp. 73–99). Springer International Publishing.