

# Enhancing Explainability and Performance in Intrusion Detection Systems using Deep Learning Models and LLMs

UNIVERSITY OF TURKU  
Department of Computing  
Master of Science (Tech) Thesis  
Cyber Security Engineering  
July 2025  
Mohd Junaid Ahmed

Supervisors:  
Tahir Mohammad (University of Turku)  
Ismayil Hasanov (University of Turku)

UNIVERSITY OF TURKU  
Department of Computing

MOHD JUNAID AHMED: Enhancing Explainability and Performance in Intrusion  
Detection Systems using Deep Learning Models and LLMs

Master of Science (Tech) Thesis, 80 p., 14 app. p.  
Cyber Security Engineering  
July 2025

---

The evolving landscape of cyber security highlights the importance of the effectiveness and transparency of Intrusion Detection Systems (IDS), which play a critical role in protecting computer networks from malicious activities. However, many advanced Machine Learning (ML) models used in IDS often present challenges in interpretation, thereby constraining their reliability and practical deployment. This research aims to improve the detection performance and explainability of IDS by combining powerful tabular Deep Learning (DL) models with open-source Large Language Models (LLMs).

In this study, the CSE-CIC-IDS2018 dataset is used as a benchmark for training and evaluating several ML models. The models include TabNet, a DL model specifically designed for tabular data, and various AutoGluon-based models such as Neural Network implemented with PyTorch (NN\_TORCH), Gradient Boosting Machine (GBM), Categorical Boosting (CaTBoost), and Extreme Gradient Boosting (XGBoost). These models are evaluated based on their ability to detect different kinds of network intrusions reliably.

After making predictions, the outputs obtained from the models are passed to open-source LLMs which generate natural language explanations. This step is intended to make the decision-making process of the models more understandable to human, including security analysts and system administrators. By integrating LLMs into the IDS pipeline, the system not only effectively identifies threats but also provides an explanation of the reasons behind each prediction in a human-readable format. The experimental findings indicate that the suggested method delivers a high detection accuracy. AutoGluon ensembles attained up to 98.1% accuracy, while TabNet achieved 97.8%. Additionally, this approach provides clear and beneficial explanations through LLMs. Although the performance wasn't consistent across all minor attack classes, integrating DL with LLMs significantly increased the system's transparency and utility for analysts. This improvement in understandability enhances the system's practical applicability in cyber security contexts.

Keywords: Intrusion Detection Systems, Tabular Deep Learning, TabNet, AutoGluon, Neural Networks, Gradient Boosting, Explainable AI, Large Language Models, CSE-CIC-IDS2018, Cyber Security

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Background . . . . .	1
1.2	Problem Statement . . . . .	2
1.3	Research and Development Need . . . . .	3
1.4	Research Questions and Objectives . . . . .	4
1.4.1	Research Questions . . . . .	4
1.4.2	Objectives . . . . .	4
1.5	Contributions . . . . .	5
1.6	Structure of the Thesis . . . . .	5
1.7	Statement of AI Usage . . . . .	6
<b>2</b>	<b>Background and Related Work</b>	<b>8</b>
2.0.1	Key AI Concepts: . . . . .	8
2.0.2	Additional Concepts and Evaluation Metrics in AI-driven IDS	9
2.1	Intrusion Detection Systems (IDS) . . . . .	11
2.1.1	Evolution of IDS Approaches: From Rule-Based to LLM- Integrated Systems . . . . .	12
2.2	Explainability in Machine Learning . . . . .	13
2.2.1	Importance in Cyber Security . . . . .	13
2.2.2	SHAP and LIME . . . . .	14

2.2.3	LLM-Based Explanations . . . . .	14
2.2.4	Operational Integration in SOCs . . . . .	15
2.3	Tabular Deep Learning Models . . . . .	16
2.3.1	TabNet . . . . .	17
2.3.2	AutoGluon: Automated ML for Tabular Data . . . . .	17
2.4	Large Language Models (LLMs) for Explanation . . . . .	18
2.4.1	Open-Source LLMs . . . . .	18
2.5	Relevance to Security Decision Support . . . . .	20
2.6	Related Work . . . . .	21
2.6.1	Deep Learning in IDS . . . . .	21
2.6.2	XAI and LLM Integration . . . . .	21
2.6.3	LLMs in IDS Pipelines . . . . .	22
2.6.4	Comparative Summary of Related Studies . . . . .	22
2.6.5	Gaps in Existing Research . . . . .	22
2.6.6	Challenges in LLM-Based IDS Research . . . . .	24
2.6.7	Summary of Research . . . . .	25
<b>3</b>	<b>Description of Dataset and Preprocessing</b>	<b>26</b>
3.1	Dataset Description . . . . .	26
3.1.1	Overview . . . . .	26
3.1.2	Feature Structure . . . . .	27
3.1.3	Attack Types . . . . .	28
3.1.4	Original Data Source and Crediting . . . . .	29
3.2	Data Cleaning and Feature Engineering . . . . .	29
3.2.1	Handling Missing Values . . . . .	29
3.2.2	Feature Selection . . . . .	30
3.2.3	Encoding Categorical Features . . . . .	31
3.2.4	Normalization and Scaling . . . . .	31

3.3	Data Splitting . . . . .	31
3.3.1	Splitting Strategy . . . . .	32
3.3.2	Rationale for Stratification . . . . .	32
3.3.3	Temporal Decorrelation . . . . .	33
3.4	Justification of Dataset . . . . .	33
3.4.1	Realistic and Diverse Attack Scenarios . . . . .	33
3.4.2	Rich Feature Set for Deep Learning Models . . . . .	33
3.4.3	Explainability Relevant Characteristics . . . . .	34
3.4.4	Standard Benchmark for Reproducibility . . . . .	34
<b>4</b>	<b>Design and Specification of the Proposed System</b>	<b>35</b>
4.1	System Architecture Overview . . . . .	35
4.2	Model Selection Rationale . . . . .	36
4.2.1	Why TabNet? . . . . .	36
4.2.2	Why AutoGluon? . . . . .	37
4.2.3	Comparative Strengths . . . . .	37
4.3	Design Considerations . . . . .	38
4.3.1	Trade-offs Between Performance and Interpretability . . . . .	38
4.3.2	Real-Time Usability Considerations . . . . .	38
4.3.3	Usability for Security Analysts . . . . .	39
4.4	Tools and Platforms . . . . .	40
4.4.1	Software Frameworks and Libraries . . . . .	40
4.4.2	Hardware and Execution Environments . . . . .	41
4.4.3	Training Efficiency and Resource Considerations . . . . .	41
4.4.4	Reproducibility and Open Source Tools . . . . .	42
4.5	LLM Integration for Explainability . . . . .	43
4.5.1	System Design for LLM Integration . . . . .	43
4.6	Choice of Open-Source LLMs . . . . .	43

4.6.1	Prompt Engineering Strategies . . . . .	44
4.6.2	Handling Failures and Output Storage . . . . .	44
4.7	Evaluation Criteria . . . . .	45
4.7.1	Performance Evaluation Metrics . . . . .	45
4.7.2	Explainability Evaluation Metrics . . . . .	46
<b>5</b>	<b>Implementation and Verification</b>	<b>47</b>
5.1	Model Training and Tuning . . . . .	47
5.1.1	TabNet Training Configuration . . . . .	47
5.1.2	AutoGluon Training Configuration . . . . .	48
5.2	Performance Evaluation . . . . .	49
5.2.1	AutoGluon Model Results . . . . .	50
5.2.2	TabNet Results . . . . .	51
5.2.3	Comparative Analysis: TabNet vs AutoGluon . . . . .	52
5.2.4	Per-Class Performance Analysis . . . . .	53
5.3	LLM Explanation Outputs . . . . .	54
5.3.1	Purpose of LLM Integration . . . . .	55
5.3.2	Models Used and Prompting Strategy . . . . .	55
5.3.3	Sample Explanation Output . . . . .	55
5.4	Validation Techniques . . . . .	56
5.4.1	Cross-Validation . . . . .	57
5.4.2	Confusion Matrix . . . . .	57
5.5	Key Findings . . . . .	59
5.6	Results in Relation to Research Questions . . . . .	59
5.7	Error Analysis . . . . .	60
<b>6</b>	<b>Discussion</b>	<b>62</b>
6.1	Performance Analysis . . . . .	62

6.1.1	Insights from Model Performance Results . . . . .	63
6.1.2	Trade-offs Between Models . . . . .	63
6.2	Explainability Evaluation . . . . .	64
6.2.1	Qualitative Evaluation of LLM Output . . . . .	64
6.3	Analyst Usability Discussion . . . . .	66
6.4	Limitations . . . . .	68
6.4.1	Dataset Coverage and Scalability . . . . .	68
6.4.2	Limits of LLM-Generated Explanations . . . . .	69
6.5	Implications for Cyber security . . . . .	70
6.5.1	Operational Integration into SOC Workflows . . . . .	70
6.5.2	Human-in-the-Loop Threat Analysis . . . . .	71
6.5.3	Compliance and Reporting Support . . . . .	71
6.5.4	Future Directions in Cyber Defense . . . . .	71
<b>7</b>	<b>Conclusion and Future Work</b>	<b>73</b>
7.1	Summary of Contributions . . . . .	73
7.2	Main Findings . . . . .	74
7.2.1	Performance Uplift via Tabular DL Models . . . . .	74
7.2.2	Enhanced Interpretability via LLMs . . . . .	75
7.2.3	Robust Evaluation Metrics Support Findings . . . . .	75
7.2.4	Class-wise Observation . . . . .	76
7.3	Impact Reflection . . . . .	76
7.3.1	Bridging the Performance–Explainability Gap in IDS . . . . .	77
7.3.2	Elevating Human-in-the-Loop Analysis . . . . .	77
7.4	Future Work . . . . .	77
	<b>References</b>	<b>81</b>
	<b>Appendices</b>	

<b>A</b>	<b>Sample Code Snippets</b>	<b>A-1</b>
A.1	TabNet Configuration . . . . .	A-1
A.2	AutoGluon Training Setup . . . . .	A-1
<b>B</b>	<b>LLM Prompt-Responses</b>	<b>B-3</b>
B.1	Prompt Template . . . . .	B-3
B.2	Response from Mistral-7B-Instruct . . . . .	B-5
B.3	Response from DeepSeek-R1 . . . . .	B-6
B.4	Response from LLaMA3-8B-Instruct . . . . .	B-9
B.5	Response from Mixtral-8x7B-Instruct . . . . .	B-11
B.6	Response from Qwen2.5 . . . . .	B-12
B.7	Response from Qwen3 . . . . .	B-13

# List of Figures

4.1	System architecture for IDS with LLM explanation. Created using draw.io. . . . .	36
5.1	Per-class performance metrics showing Precision, Recall, and F1-Score for each network traffic class. The plot highlights strong performance for most classes but low scores for Infiltration and Web Attack.	52
5.2	Comparison of key metrics (Accuracy, Precision, F1-Score) between TabNet and AutoGluon. . . . .	54
5.3	LLM Explanation Word Counts across Models . . . . .	57
5.4	TabNet Confusion Matrix . . . . .	58
5.5	AutoGluon Confusion Matrix . . . . .	58
6.1	Flesch-Kincaid Grade Level (Lower is Easier) . . . . .	67
6.2	Flesch Reading Ease Score (Higher is Easier) . . . . .	68
7.1	System architecture sketch showing real-time streaming IDS with async LLM explanation flow. . . . .	78
7.2	Flowchart showing SHAP-LLM hybrid explanation pipeline. . . . .	79

# List of Tables

2.1	Comparison of studies using DL models and LLM/XAI methods for IDS interpretability . . . . .	23
3.1	Distribution of Network Traffic by Label . . . . .	28
4.1	Software frameworks for model training and evaluation . . . . .	40
4.2	LLM-related libraries used for inference and backend integration . . .	40
4.3	Visualization libraries used in the IDS and LLM explanation pipeline	41
4.4	Execution hardware and runtime environment on Google Colab . . .	41
4.5	Summary of platforms, models, frameworks, and system specifications	42
4.6	Comparison of Instruction-Tuned Language Models . . . . .	45
4.7	Performance Metrics Used . . . . .	45
4.8	Explainability Metrics . . . . .	46
5.1	Summary of TabNet training configuration including core hyperparameters and training settings. . . . .	48
5.2	Summary of AutoGluon training configuration including model types, ensemble strategy, and best model selection. . . . .	49
5.3	Other Notable Models . . . . .	51
5.4	Per-Class Performance of TabNet . . . . .	52
5.5	Comparison of TabNet and AutoGluon (Best Model) . . . . .	53
5.6	Overview of Selected LLMs Used . . . . .	55

6.1	Flesch-Kincaid Grade Level interpretation. . . . .	66
6.2	Readability Metrics for LLM Outputs (Measured via Flesch-Kincaid Calculator) . . . . .	67
7.1	Comparison of TabNet and AutoGluon (Best Model) [4], [5] . . . . .	75
7.2	Per-Class Performance of TabNet . . . . .	76
7.3	Dataset comparison matrix. . . . .	79

# List of acronyms

**API** Application Programming Interface

**APIs** Application Programming Interfaces

**AutoGluon** AutoML toolkit developed by AWS for structured/tabular data

**CaTBoost** Categorical Boosting

**CIC** Canadian Institute for Cyber security

**CNN** Convolutional Neural Network

**CSE-CIC-IDS2018** Canadian Institute for Cyber Security – Intrusion Detection System 2018 Dataset

**DL** deep learning

**GBM** Gradient Boosting Machine

**IAT** Inter-Arrival Time

**IDS** Intrusion Detection Systems

**LIME** Local Interpretable Model-Agnostic Explanations

**LLMs** Large Language Models

**ML** machine learning

**NLP** Natural Language Processing

**NN\_TORCH** Neural Network implemented with PyTorch (within AutoGluon)

**RNN** Recurrent Neural Network

**SHAP** SHapley Additive exPlanations

**SOCs** Security Operations Centers

**TabNet** Tabular Network, a deep learning architecture for tabular data

**XAI** Explainable AI

**XGBoost** Extreme Gradient Boosting

# 1 Introduction

The complexity of modern cyber threats necessitates a more advanced approach than traditional signature-based methods. These methods have become inadequate against adaptive and evolving attacks. In this thesis, a hybrid framework is proposed to improve the efficiency of IDS.

Deep Learning (DL) has shown promise in enhancing detection capabilities, but its decision-making processes can be opaque, reducing trust and potentially delaying incident responses. To address this issue, the proposed framework combines advanced tabular DL models with explanations from LLMs.

The goal of this approach is to improve both the accuracy of detection in IDS and to improve the explainability of its decisions, making it easier for humans to understand and trust the system's outputs. By integrating the strengths of both DL and LLMs, the hybrid framework aims to provide a more robust and transparent solution for IDS.

## 1.1 Background

Cyber security has become vital for digital infrastructure, since networks encounter more complex threats. Intrusion Detection Systems (IDS) are now playing a key role in defending from malicious activities and unauthorized access. IDS monitors network traffic, identifies potential threats, and alert security professionals in real-time. IDS approaches can be classified in broad terms. They do include signature-

based systems as well as anomaly-based systems. Anomaly-based systems utilize ML to identify deviations beyond normal behavior while signature-based systems rely on predefined attack patterns [1].

Advances in Machine Learning (ML), particularly DL, have resulted in the creation of precise models for anomaly detection. However, these advanced models, including deep neural networks and ensemble methods, often function as black boxes, making their underlying decision-making processes unclear. This lack of transparency poses substantial issues in critical domains such as cyber security, where understanding the rationale behind a detection is crucial for trust and swift response [2].

The structure of cyber security datasets, like CSE-CIC-IDS2018 [3], is fundamentally tabular and includes both categorical and numerical features with imbalanced classes. This poses difficulties for standard DL architectures that are typically designed to process unstructured data such as images and text. Advancements in tabular DL models, such as TabNet [4], and automated ML frameworks like AutoGluon [5] have shown potential for handling tabular data with competitive performance. However, their application in IDS requires solutions that address the dual challenges of high performance and explainability.

## 1.2 Problem Statement

DL models have greatly enhanced the precision of IDS. However, their lack of transparency poses a challenge for their implementation in practical scenarios. Cyber security analysts need not just accurate alerts but also an understanding of why a particular traffic pattern was identified as malicious. Without interpretability, models cannot be fully trusted, which can delay incident responses or lead to misinterpretation of benign events as threats (false positives) [6].

Traditional explainability methods such as SHAP [7] or LIME provide some

insights into model behavior, but they often produce technical explanations that may not be understandable to all stakeholders. Moreover, they do not generate context-aware, natural language explanations that could facilitate quicker decision-making.

Additionally, there is a lack of approaches that integrate both high-performance detection models and human-interpretable explanations, especially ones tailored to tabular cyber security datasets. This creates a performance-explainability tradeoff, where enhancing one often degrades the other. The gap between predictive performance and interpretability in IDS remains a critical shortcoming in current research and applications.

### 1.3 Research and Development Need

To address these challenges, there is a clear need for a hybrid IDS framework that not only achieves high detection accuracy using tabular DL models but also integrates explainability mechanisms capable of producing understandable and actionable insights. The recent progress in open-source LLMs, such as GPT-J, Falcon, and LLaMA, presents an opportunity to translate complex model outputs into natural language explanations, which are more accessible to human users [8].

Combining TabNet and AutoGluon-based model including NN\_TORCH, GBM, CaTBoost and XGBoost with LLM-driven explanation modules can provide a robust and interpretable IDS pipeline. Such a system can potentially support analysts by making DL models' reasoning transparent and improving response effectiveness. Moreover, using open-source LLMs ensures transparency, reproducibility, and adaptability to various environments without depending on proprietary APIs.

## 1.4 Research Questions and Objectives

The focus of this thesis is on the discrepancy between advanced IDS models and their explainability. While models such as TabNet and AutoGluon exhibit high precision in detecting threats, their decision-making processes are typically not transparent, which is crucial for trust and efficient real-time security operations. Consequently, this research aims to determine if natural language explanations generated by LLMs can help bridge this gap.

### 1.4.1 Research Questions

- **RQ1:** How do tabular deep learning models like TabNet and AutoGluon perform on a cyber security dataset like CSE-CIC-IDS2018?
- **RQ2:** Can LLMs effectively generate accurate and understandable explanations for predictions made by tabular models?
- **RQ3:** What is the trade-off between performance and explainability when integrating LLMs in the IDS pipeline?

### 1.4.2 Objectives

- To evaluate the predictive performance of TabNet and AutoGluon (including NN\_TORCH, GBM, CaTBoost, and XGBoost) on the CSE-CIC-IDS2018 dataset.
- To design and implement a modular IDS pipeline that integrates model outputs with open-source LLMs for explanation.
- To develop a reproducible, open-source solution for explainable IDS frameworks.

## 1.5 Contributions

In this thesis, a novel IDS pipeline is proposed that integrates high-performance tabular DL models, namely TabNet and AutoGluon ensembles (NN\_TORCH, GBM, CaTBoost, XGBoost) with open-source LLMs to generate human-readable explanations of detection outcomes. This approach addresses the dual challenges of explainability and performance in IDS, using the CSE-CIC-IDS2018 dataset as a benchmark.

The core contributions of this thesis include:

- A performance evaluation of various tabular DL models on the CSE-CIC-IDS2018 dataset.
- A modular framework for integrating LLMs to interpret model outputs in natural language.
- A qualitative and quantitative analysis of the generated explanations to assess trustworthiness and interpretability.
- An open-source implementation pipeline that supports reproducible research in explainable cyber security systems.

This contribution supports cyber security analysts in understanding model decisions more clearly, potentially increasing their trust in automated alerts, reducing investigation time, and enabling more informed incident response actions.

## 1.6 Structure of the Thesis

The rest of the thesis is organized as follows:

*Chapter 2* provides an overview of the theoretical background, discussing the principles of IDS, the significance of explainability in ML, and the utilization of

tabular DL models. Particular attention is given to LLMs for generating understandable explanations. The chapter ends with a review of relevant literature and its connection to security decision-making.

*Chapter 3* explains the dataset and data preparation methods, such as data cleaning, feature engineering, encoding, normalization, and data-splitting strategies. The selection of the dataset is justified in accordance with the research objectives.

*Chapter 4* outlines the design of the proposed IDS framework, including the system architecture, reasons for model selection, and design considerations. It also includes information on the integration of LLMs for explainability and evaluation criteria for both performance and interpretability.

*Chapter 5* describes the implementation and verification process, encompassing model training and adjustment, experimental setup, performance evaluation, and LLM-generated explanations. Techniques used to guarantee result dependability are also detailed.

*Chapter 6* discusses the results, examining performance, explainability, and usability. The limitations of the approach and its implications for cyber security operations are analyzed.

*Chapter 7* summarizes the contributions and findings, assesses the research impact, and suggests future work to enhance the framework's applicability in real-world security contexts.

## 1.7 Statement of AI Usage

*AI tools were used for searching and gathering the research papers. Tools such as Elicit and Google Scholar were used to discover academic literatures. These tools were helpful because they allowed me in citations, creating interactive mind maps of research topics, and group papers into collections. While AI was employed to refine the written content and eliminate any grammatical errors and writing flaws, it is*

*important to note that the content was exclusively studied, analysed, researched and written by me.*

## 2 Background and Related Work

In the context of escalating cyber threats, there is an increasing need for IDS to be more precise, understandable, and swift. Conventional IDS methods frequently lack the necessary adaptability to manage new attacks. DL has proven beneficial in enhancing IDS precision, but its non-transparent decision-making process poses a challenge in critical applications. A promising solution is the combination of tabular DL models with LLMs, which offers both high performance and interpretable explanations. This section aims to discuss fundamental concepts in IDS, explainability in ML, tabular DL techniques, and recent advancements in IDS systems integrating LLMs.

### 2.0.1 Key AI Concepts:

This section provides an overview of fundamental AI and ML principles that are crucial for this study. It explains basic techniques, models, and strategies that form the basis of the proposed hybrid IDS. These concepts are vital in understanding how the system identifies threats, processes tabular data, and generates easily understandable explanations for humans.

- **Machine Learning (ML):** Algorithms that learn patterns from data, enabling them to make predictions or decisions without being explicitly programmed.

- **Deep Learning (DL):** A subset of ML that uses multi-layered neural networks to model complex relationships in the data.
- **Transfer Learning:** The use of a pre-trained model for a new task, saving time and resources by leveraging existing knowledge.
- **Automated Machine Learning (AutoML):** Automates the process of selecting and tuning models for improved efficiency.
- **Explainable AI (XAI):** Methods that make AI model decisions comprehensible to humans, ensuring transparency and trust.
- **SHAP:** A model-agnostic method used to explain predictions by attributing outputs to specific input features, based on Shapley values from game theory.
- **LIME:** A model-agnostic method that generates interpretable local approximations of model predictions.
- **TabNet:** A deep learning architecture designed specifically for tabular data, using attention mechanisms to prioritize relevant features.
- **AutoGluon:** An open-source AutoML toolkit focused on structured/tabular data, providing pre-trained models and algorithms for easy implementation.
- **Large Language Models (LLMs):** AI systems trained to process and generate human language, enabling them to interact with users in a natural, conversational manner.

## 2.0.2 Additional Concepts and Evaluation Metrics in AI-driven IDS

While core AI concepts form the theoretical basis, understanding additional technical and operational notions is crucial for designing, evaluating, and ensuring the

reliable performance of AI-driven IDS in real-world scenarios. These supplementary concepts include model transparency, evaluation metrics, and cyber security-specific terms. This section outlines key ideas related to interpreting results, avoiding common mistakes, and maintaining the IDS's effectiveness in actual operational settings.

- **Black Box Model:** A model whose inner workings are not clear and hard to understand by humans, making it challenging to explain why specific predictions are made.
- **White Box Model:** A model whose internal mechanisms can be understood, enabling analysts to comprehend how inputs are transformed into outputs.
- **Feature Engineering:** The process of selecting, altering, or creating input variables (features) to improve ML models performance
- **Class Imbalance:** A frequent problem in cyber security data sets where one class (e.g., regular traffic) has significantly more occurrences than another (e.g., rare attack types), which can lead to biased models toward the majority classes.
- **Precision:** The percentage of true positives among all predicted positives.
- **Recall:** The proportion of actual positives among all identified positives.
- **F1-Score:** The average of precision and recall, balancing both metrics.
- **Confusion Matrix:** A table used to evaluate classification model performance, displaying true positives, false positives, true negatives, and false negatives.
- **Overfitting:** When a model learns patterns specific to the training data (including noise) but fails to generalize to unseen data.

- **Underfitting:** When a model is too simple to capture the underlying patterns in the data, resulting in poor performance on both training and testing datasets.
- **Data Leakage:** When information from outside the training dataset (e.g., test set) is accidentally used during model training, causing overly optimistic performance estimates
- **Adversarial Example:** Input data intentionally designed to cause a machine learning model to make an incorrect prediction.
- **Latency:** In IDS, the time taken for the system to detect and respond to a potential threat after the data is observed
- **Security Operations Center (SOC):** A centralized team responsible for monitoring, detecting, investigating, and responding to cybersecurity threats in real-time.
- **Zero-Day Attack:** A cyber attack that targets an unknown vulnerability, making it difficult for traditional signature-based IDS to detect
- **Benchmark Dataset:** A standard dataset used for comparing model performance across various research studies.

## 2.1 Intrusion Detection Systems (IDS)

IDS play a vital role in protecting network infrastructure from unauthorized access and malicious activities. These systems monitor and analyze network traffic to identify potential intrusions or unusual patterns, alerting administrators for prompt action. IDS techniques can be divided into two main categories: signature-based and anomaly-based detection systems [9].

In simpler terms, signature-based IDS utilize predefined patterns of malicious actions, similar to antivirus databases, effectively identifying known security threats. However, they struggle with detecting novel or zero-day attacks because these attacks lack the familiar patterns. On the contrary, anomaly-based IDS create a standard of normal network behavior and flag deviations from this norm as potential threats. Anomaly-based systems provide enhanced adaptability to new types of attacks; however, they frequently experience high false-positive rates and difficulties in interpretation.

The escalating complexity of cyber-attacks necessitates the use of ML and DL methods in anomaly-based IDS. However, many DL-based IDS are still opaque, making it hard for cyber security experts to understand their decisions. Recent studies have emphasized this challenge and suggest combining XAI techniques with LLMs to make these systems more transparent and useful for cyber security professionals [10], [11].

### **2.1.1 Evolution of IDS Approaches: From Rule-Based to LLM-Integrated Systems**

In the early stages of IDS, rule-based and signature-based methods such as Snort and Bro/Zeek were predominantly used. These systems compared network activity against known attack patterns, performing well for established threats but failing to detect new or evolving attacks [12], [13].

In the 2000s, traditional ML methods such as Decision Trees, Support Vector Machines (SVMs), and Random Forests were applied to anomaly detection tasks. While these models improved generalization, they still required manually engineered features and lacked interpretability for non-expert users [14], [15].

From the mid-2010s, DL approaches such as CNNs, RNNs and autoencoders became more effective than traditional models in terms of accuracy, particularly for

large and complex datasets like NSL-KDD and CSE-CIC-IDS2018 [16]. However, their black-box nature led to trust issues in operational settings [17].

More recently, XAI techniques such as SHAP and LIME have emerged to help security analysts understand the reasons behind triggered alerts. These tools enhance transparency but often provide explanations that are too technical or disconnected from real-world analyst workflows [12].

The latest trend is the integration of LLMs with IDS. LLMs can generate natural language justifications for model predictions, bridging the gap between high accuracy and human interpretability. This represents a significant stage in IDS evolution, where automated detection is coupled with human-understandable reasoning [18].

## 2.2 Explainability in Machine Learning

In cyber security applications, explainability is not a luxury, it is a necessity. Analysts must understand why an IDS flags certain behaviors as threats, especially in high-stakes environments such as banking, defense, or healthcare. Without interpretability, even high-performing models can lead to operational friction, mistrust, or misinformed decisions. As argued by Doshi-Velez and Kim [2], explainability is essential for human trust, safety, and compliance.

### 2.2.1 Importance in Cyber Security

Cyber security analysts often rely on IDS outputs for real-time decision-making. If a system raises an alert but cannot explain why, it may be ignored or misinterpreted, especially in cases of false positives. Explainability enables the identification of model biases, improves debugging of false alarms, and enhances collaboration between human analysts and automated systems. Moreover, with regulatory frameworks like the EU AI Act [19] pushing for algorithmic transparency, explainability

is becoming a compliance requirement.

### 2.2.2 SHAP and LIME

Two prominent model-agnostic explainability tools are **SHAP** (**SH**apley **A**dditive **e**x**P**lanations) and **LIME** (**L**ocal **I**nterpretable **M**odel-**A**gnostic **E**xplanations).

- **SHAP** provides consistent and theoretically grounded explanations by assigning feature importance scores based on cooperative game theory. It has been widely used in cyber security contexts to explain anomaly detections in ML models [7].
- **LIME** perturbs the input data and fits local surrogate models to approximate the decision boundary of the black-box classifier. Although faster, LIME can be unstable and sensitive to sampling variations [20].

Both SHAP and LIME, however, generate technical outputs like bar graphs or feature attributions, which are not always intuitive to non-ML experts. Furthermore, they do not offer contextualized explanations or rationale in natural language.

### 2.2.3 LLM-Based Explanations

In the advancement of LLMs, such as LLaMA, Mistral, and Qwen3, there has been a significant shift in the approach to explainability. Rather than using complex diagrams or numeric feature weights, these models can convert IDS outputs into understandable justifications that are similar to analyst notes or reports. For example, a LLM could interpret a prediction labeled as "DDoS Attack" and provide an explanation such as: "This traffic pattern is similar to known DDoS behavior due to its high packet rate, frequent targeting of the same destination, and low entropy."

The most recent studies, such as Houssel et al. [21] and Ali [22], indicate that integrating LLMs into IDS can be advantageous. This integration offers improved interpretability without compromising performance. By doing so, these systems can offer real-time explanations in a narrative format for both accurate and inaccurate predictions. This, in turn, boosts the effectiveness of incident response.

Researchers have begun to combine SHAP outputs with LLM prompts. This integration allows for the delivery of both numerical and descriptive explanations. This approach aims to overcome the limitations of traditional tools [23].

#### 2.2.4 Operational Integration in SOCs

Integrating LLM based explanations into SOC workflows can significantly improve incident handling speed and accuracy. Here's a simplified explanation:

- **Alert Prioritization:** LLM generated explanations help analysts swiftly determine if a flagged event is a high-severity incident or a false positive, thereby reducing time spent on low-priority alerts. As demonstrated by Houssel et al. [21], severity-ranked explanations enhance SOC triage efficiency.
- **Context-Enriched Incident Reports:** Instead of numerical outputs, LLMs can generate easy-to-understand summaries like "The model classified this traffic as an SQL injection attempt due to abnormal query patterns, frequent malformed SQL keywords, and repeated connection resets to the database server." This is in line with findings from Yang et al. [23], where narrative explanations improved communication between SOC analysts and non-technical stakeholders.
- **Training and Knowledge Retention:** SOCs can use LLM generated explanations as a growing knowledge base. By storing explanations alongside

alerts, analysts can create historical case libraries that speed up onboarding for new staff, as discussed in Wali et al [24].

In essence, LLM based explanations are not just academic tools but also operational boosters for real-time cyber defense. This supports the argument for hybrid IDS frameworks that combine high-performance Deep Learning models with advanced, human-readable explainability mechanisms.

## 2.3 Tabular Deep Learning Models

DL has demonstrated exceptional results in domains handling unstructured data like images and text. However, in the context of cyber security applications such as IDS, data is usually structured in a tabular format, consisting of numerical and categorical features that describe various characteristics of network traffic. Traditional DL architectures are not optimized for this structure, often necessitating extensive feature engineering or transformations which can affect the model's transparency and performance.

To overcome these challenges, recent developments have introduced specialized tabular DL models capable of directly processing structured data. This thesis focuses on two notable methods: TabNet, which utilizes attention mechanisms for feature selection and interpretability, and AutoGluon, An Automated Machine Learning (AutoML) framework that incorporates a wide range of high-performing models like NN\_TORCH, GBM, CaTBoost, and XGBoost. These models have displayed impressive performance on cyber security datasets and serve as a base for integrating explainability techniques through model-agnostic tools and LLMs.

### 2.3.1 TabNet

IDS often rely on tabular data structured as records containing numerical and categorical features like packet size, protocol, source/destination IPs, and flow duration. While traditional ML models (e.g., Random Forest, SVM) perform well on such data, tabular DL has shown competitive or superior performance when properly optimized. However, handling the complexity of tabular structures such as feature interactions, sparsity, and imbalance remains a challenge.

**TabNet**, introduced by Arik and Pfister [4], is a DL model tailored for tabular datasets. It employs sequential attention mechanisms to select and focus on relevant features during each decision step. This provides both high predictive performance and built-in interpretability by showing which features were used at which step of the decision process.

TabNet's ability to select important features dynamically is crucial, especially when handling high-dimensional network traffic data with mixed-type attributes.

### 2.3.2 AutoGluon: Automated ML for Tabular Data

AutoGluon is an automated ML(AutoML) framework created by AWS that optimally combines various models to achieve the best performance on tabular datasets [5]. It's particularly beneficial for IDS, as feature engineering and hyperparameter tuning can be intricate and time-consuming. The primary models utilized within AutoGluon are:

- **NN\_TORCH**: This is a neural network built using PyTorch, which can be customized with different layers and activations. It's suitable for handling large datasets, even those with noise, and it leverages GPU acceleration to boost processing speed.
- **GBM**: This is a popular ensemble learning technique that creates models

sequentially to minimize residual errors. It performs well on moderately imbalanced data, particularly in cyber security datasets.

- **CaTBoost:** Developed by Yandex, this algorithm can handle categorical features directly and prevents overfitting using ordered boosting. Its fast convergence and high accuracy make it ideal for structured cyber security data.
- **XGBoost:** This is a powerful ML tool recognized for its scalability and regularization. It's widely used in IDS due to its exceptional performance across binary and multiclass classification tasks.

The latest studies on IIDS suggest that ensemble methods, like AutoGluon, generally outperform individual models. This is because ensemble methods provide model diversity and optimized pipelines.

Furthermore, AutoGluon can easily be combined with explainability tools, making it suitable for hybrid explainable IDS pipelines when used alongside SHAP or LLMs [24].

## 2.4 Large Language Models (LLMs) for Explanation

The advancement of LLMs marks a substantial progress in ML systems capacity to explain their thought processes to human users. In the domain of cyber security, where clear and swift explanations are vital, LLMs offer a revolutionary approach for translating complex model outcomes into easy-to-understand narratives.

### 2.4.1 Open-Source LLMs

In the field of cyber security, advancements have been made in the utilization of open-source LLMs within IDS. These innovative models can generate contextually

relevant, easy-to-understand explanations for predictions made by the model, converting complex outputs into digestible insights for analysts. This research evaluates various state-of-the-art open-source LLMs to ascertain their potential in improving IDS interpretability and user-friendliness, with a focus on real-time applications.

- **The Mistral-7B** model, developed by Mistral AI, is an optimized transformer model designed for real-time applications such as intrusion alert summarization. It has been fine-tuned to provide context-aware justifications, making it useful for analysts [22].
- **DeepSeek-R1** is a language learning model that integrates language and coding capabilities. It supports deep packet inspection and offers coherent explanations for alerts in critical infrastructure systems [25].
- **The LLaMA3-8B-Instruct** model, released by Meta, performs well at a moderate scale. It is particularly beneficial for multilingual intrusion detection system interpretation and threat sequence summarization [26].
- **Mixtral-8x7B** model is the mixture-of-experts model, optimized for efficiency, is designed for large-scale explanation tasks in Security Operations Center (SOC) environments. It integrates seamlessly with IDS dashboards [10].
- **Qwen2.5-14B-Instruct** is recognized for its capability to follow instructions precisely and multilingual support. It converts complex feature importance, such as SHAP outputs, into straightforward and natural language explanations [27].
- **Qwen3-8B-Base** is a base model, it offers advanced prompt engineering for log analysis and threat chain explanation. This model can be fine-tuned to enhance its cyber reasoning skills [28].

All three models can be fine-tuned or prompted to provide context-aware justifications for IDS decisions, such as why a packet was classified as malicious or what features triggered an anomaly alert. For example, Houssel et al. [21] demonstrated that LLMs, when paired with model outputs from classifiers, can generate natural language justifications with high coherence and utility for human analysts.

## 2.5 Relevance to Security Decision Support

The key advantage of LLMs in IDS is their ability to act as explanation generators that are both technically informative and linguistically accessible. For instance, instead of showing a feature attribution chart, an LLM can generate a sentence like:

`“The model identified this traffic as part of a brute-force attack due to high-frequency login attempts, unusual access times, and repeated password failures.”`

This form of explanation aligns better with the needs of security operations teams, who often need to act quickly and may not be trained in data science. LLMs allow for interactive querying, meaning analysts can ask questions like “Why was this IP flagged?” or “What type of attack does this resemble?”, creating a conversational interface with machine intelligence.

Recent studies, such as the one by Houssel et al. [21], demonstrate that integrating LLMs with traditional IDS enhances analyst response time and clarifies incidents. Additionally, the study conducted by Yang et al. [23] evaluated several LLMs using intrusion datasets. They found that the models consistently provided high-quality explanations, were fluent, and relevant across different categories of cyber attacks.

## 2.6 Related Work

The focus of research is broadening to enhance IDS by integrating DL and explainable AI (XAI) techniques, such as LLMs. This synergistic approach aims to accomplish three main goals: improving detection accuracy, guaranteeing transparency in decision-making procedures, and enabling quick decisions in high-risk situations.

### 2.6.1 Deep Learning in IDS

In the realm of IDS, conventional approaches like Decision Trees and Random Forests are being progressively augmented or even replaced by DL models, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), TabNet, and transformers. As opposed to traditional methods, DL models exhibit superior pattern recognition abilities, especially in complex datasets like CSE-CIC-IDS2018 or NSL-KDD.

In their work, Alnahdi & Narain [29] introduced a coherence-based DL framework. This unique framework utilizes LLMs to verify the predictions made by IDS. Compared to SHAP or LIME used separately, this model showed improved interpretability and consistency.

Similarly, T. Ali [22] presented a real-time IDS system that utilizes DL models for anomaly detection and LLMs for generating adaptive data summaries. Their evaluation on CICIDS2017 and CICIDS2018 datasets demonstrated state-of-the-art performance in terms of accuracy and interpretability.

### 2.6.2 XAI and LLM Integration

In contemporary IDS research, there is a rising emphasis on integrating SHAP or LIME for feature attribution along with natural language explanations. This combined approach enhances the credibility of alerts and enables analysts to verify and

contest model decisions more efficiently. A comparative study by Yang et al. [23] on various LLMs (LLaMA, GPT-J, Falcon) for IDS showed that while all models boosted interpretability, domain-tuned LLMs provided better factual accuracy and fluency.

### 2.6.3 LLMs in IDS Pipelines

An innovative approach involves incorporating LLMs completely within IDS pipelines, not just for post-hoc explanations but as active reasoning components. Houssel et al. [21] illustrated this with a framework in which LLMs interpret IDS alerts and rank them based on severity and confidence, thereby assisting triage and response prioritization in Security Operations Centers (SOCs).

Similarly, Wali et al. [24] proposed an explainable IDS utilizing Random Forests supplemented by GPT-J. This system not only improved interpretability but also helped automate analyst reports during intrusion investigations.

### 2.6.4 Comparative Summary of Related Studies

Table 2.1 below compares key studies that integrate DL, XAI, and LLMs in IDS. Each study’s primary contributions, limitations, and differences compared to the current study are outlined.

### 2.6.5 Gaps in Existing Research

To advance the use of DL and XAI in IDS, it is essential to address several ongoing challenges:

1. **Limited Integration of LLMs:** Most IDS frameworks utilizing XAI currently employ post-hoc feature attribution tools such as SHAP or LIME [7],

Table 2.1: Comparison of studies using DL models and LLM/XAI methods for IDS interpretability

Study	Dataset(s)	DL Model(s)	XAI/LLM Method	Key Strengths	Limitations
Alnahdi & Narain [29]	CIC-IDS2017, CIC-IDS2018	CNN, RNN	LLM verification of IDS outputs	Improved interpretability using LLMs	Limited scalability; no real-time evaluation
Ali [22]	CIC-IDS2017, CIC-IDS2018	LSTM, DNN	Real-time LLM summaries	Adaptive explanations	Focus on streaming data, not rare attack detection
Yang et al. [23]	CIC-IDS2018	Various DL models	SHAP + LLMs	High factual accuracy of explanations	No standard metric for interpretability
Wali et al. [24]	CIC-IDS2018	Random Forest	GPT-J explanations	Automated analyst reports	Not optimized for DL-based IDS
Houssel et al. [21]	NSL-KDD, CIC-IDS2018	CNN, DNN	LLM-based alert ranking	Prioritizes critical alerts effectively	No integration with tabular-specific DL models
<b>This Study</b>	CSE-CIC-IDS2018	TabNet, AutoGluon	Open-source LLM explanations	Balances performance & interpretability; targets rare attack classes	Requires further work on multilingual support

[20]. Few studies have, however, explored the full integration of LLMs into IDS for generating context-aware, natural language explanations [21].

2. **Lack of Standardized Evaluation of Explanations:** Although accuracy metrics for IDS are well established, there is no widely accepted benchmark for measuring the quality of LLM-generated explanations, making it difficult to compare the interpretability of different systems [24].
3. **Performance–Explainability Trade-off:** Several studies have shown that improving model interpretability can potentially reduce predictive accuracy, and vice versa [2], [6]. Balancing both remains an open challenge in IDS research.
4. **Underrepresentation of Rare Attack Classes:** Datasets like CSE-CIC-IDS2018 are heavily imbalanced [30], [31], leading to poor model performance on rare but critical attack types such as Infiltration and Web Attacks. Few works have explored how LLM explanations can assist analysts in low-sample

scenarios.

To address these gaps, an IDS framework should integrate high-performance tabular DL models with LLM-generated explanations while providing clear evaluation metrics for both accuracy and interpretability. This thesis aims to tackle these shortcomings.

### 2.6.6 Challenges in LLM-Based IDS Research

The integration of LLMs into IDS holds potential benefits, but several obstacles persist:

- **Computational Overhead:** Large-scale LLMs necessitate substantial GPU memory and inference time. This can pose challenges in real-time SOC environments where latency is crucial [21], [32].
- **Risk of Fabricated Explanations ("Hallucinations"):** Sometimes, LLMs generate convincing but factually incorrect justifications. In a security context, this could lead analysts to make inappropriate operational decisions [23].
- **Bias and Context Sensitivity:** If the training data for LLMs contains biases, these may be perpetuated into explanations, potentially distorting an analyst's understanding of the severity or nature of a threat [22].
- **Limited Domain Specialization:** General-purpose LLMs may lack extensive cyber threat intelligence knowledge, affecting the accuracy of their explanations unless fine-tuned on domain-specific datasets [24].
- **Security and Privacy Concerns:** The use of cloud-hosted LLM APIs might spark concerns about sensitive network data being transmitted outside the organization. Open-source, on-premise deployment can help alleviate this concern but may increase hardware demands [10], [27]

Overcoming these challenges is vital for producing trustworthy, deployable LLM-enhanced IDS. Future research should focus on lightweight, domain-tuned LLMs, robust factuality checks, and hybrid pipelines that combine statistical and narrative explanations.

### 2.6.7 Summary of Research

While integrating DL models and LLMs in IDS has shown positive results for improving both accuracy and interpretability, several significant issues persist in current studies. Firstly, existing methods often fail to account for practical constraints such as limited computing resources within real-time environments. Secondly, the need for models capable of functioning effectively across various languages and cyber security domains remains unmet. Lastly, there is a lack of standardized evaluation benchmarks to assess the effectiveness of IDS explanations generated by LLMs. These deficiencies indicate that while advancements have been made, additional research is necessary to develop scalable, adaptable, and universally comparable explainable IDS.

# 3 Description of Dataset and Preprocessing

This chapter outlines the structure and characteristics of the CSE-CIC-IDS2018 dataset and details the preprocessing steps employed to clean, transform, and prepare the data for use with tabular DL models. These preprocessing steps are crucial to guarantee the quality of the data, the performance of the model, and a fair evaluation.

## 3.1 Dataset Description

The dataset utilized in this study is named CSE-CIC-IDS2018. It was developed by the Canadian Institute for Cyber security (CIC) at the University of New Brunswick and serves as a successor to the CIC-IDS2017 dataset. This dataset is one of the most extensive publicly available for evaluating IDS performance in contemporary network environments.

### 3.1.1 Overview

The CSE-CIC-IDS2018 simulation generates network traffic resembling a real enterprise network, utilizing the AWS cloud infrastructure. This simulation encompasses both harmless and malicious traffic, representing various contemporary attack types.

Over a 10-day period, 50 machines were responsible for generating this traffic, which was recorded using the CICFlowMeter tool.

The dataset contains over 80 features, categorized into several areas such as packet statistics, flow durations, Inter-Arrival Time (IAT) metrics, flag counts, and header information. The comprehensive nature of the features makes the dataset highly suitable for ML-based IDS.

### 3.1.2 Feature Structure

Each record in the dataset corresponds to a network flow and includes the following categories of features:

- **Basic flow features:** e.g., Flow Duration, Total Fwd/Bwd Packets, Bytes, IAT metrics
- **Content features:** e.g., PSH, URG, FIN flag counts
- **Time-based features:** e.g., Fwd IAT Std, Active Min, Idle Std
- **Statistical features:** e.g., Packet Length Mean, Flow Bytes/s, Packet Length Variance
- **Application-layer indicators:** e.g., Destination Port, Protocol

The dataset includes 85 columns in total, with 84 features and one label column named "Label" indicating the class (e.g., Benign, DDoS, Botnet, etc.). The distribution of the classes is presented in Table 3.1, summarized as follows:

Table 3.1: Distribution of Network Traffic by Label

Label	Count
Benign	53,287
DDoS attack	7,726
DoS attack	1,907
Botnet	1,425
Infiltration	1,204
Brute-force	970
Web attack	9

This distribution indicates a significant class imbalance, which is common in cyber security datasets and must be considered in model evaluation and sampling strategies.

### 3.1.3 Attack Types

In the field of network security, it's essential to comprehend various types and complexities of cyber attacks for developing robust detection and defense strategies. This study explores key categories of network attacks, each modeled with traffic patterns and payload structures mimicking real-world activities. By using realistic behavioral representations, the research findings are both valid and applicable

- **Botnet Activity** - refers to coordinated actions carried out by multiple compromised computers. These activities often involve spamming, DDoS attacks, or distributing malware.
- **Web-based Attacks** (e.g., SQL Injection, Cross-Site Scripting) - target web applications by exploiting vulnerabilities through malicious user input. This can result in manipulating backend databases or injecting client-side scripts.

- **Infiltration Attacks** (Internal Threats) - encompass threats originating from within a network. These may include malicious insiders or compromised internal hosts that pose a security risk to the system.
- **Brute-force Login Attempts** - involve automated attempts to access a system using numerous password combinations, typically with the intention of gaining unauthorized access.

Each attack is mapped to realistic traffic patterns and payload structures, reflecting real-world behavior.

#### 3.1.4 Original Data Source and Crediting

The dataset was curated by Iman Sharafaldin et al. [30] and published through the Canadian Institute for Cyber Security. It is publicly available and widely adopted for academic benchmarking.

## 3.2 Data Cleaning and Feature Engineering

The CSE-CIC-IDS2018 dataset, although comprehensive, requires careful preprocessing to ensure optimal performance and fairness in ML pipelines. This section outlines the key steps undertaken in data cleaning, handling missing values, feature selection, encoding, and normalization.

### 3.2.1 Handling Missing Values

The raw dataset, particularly when collected from AWS S3 or CICFlowMeter outputs, can contain inconsistencies such as missing or infinite values due to logging errors or feature extraction bugs. However, in this work, a cleaned version of the dataset was used, as prepared by the community on Kaggle [33], which ensured:

- Zero missing values
- No infinite or NaN entries
- Correctly parsed data types
- No duplicated records

Such pre-cleaned data reduces the preprocessing overhead and minimizes data integrity issues that often affect anomaly detection accuracy.

### 3.2.2 Feature Selection

From the original 85 columns, features were selected based on relevance, diversity across statistical, time-based, and flag features, and compatibility with tabular DL models. The chosen feature set of 40+ columns (listed below) excludes redundant identifiers (e.g., IP addresses, timestamps) and emphasizes those that capture flow behavior and anomaly patterns:

---

#### Listing 1 Feature Column Definition

---

```
feature_cols = [
    'Protocol', 'Fwd Packet Length Mean', 'Fwd Packet Length Std',
    'Bwd Packet Length Min', 'Flow Bytes/s', 'Flow IAT Std',
    'Fwd IAT Total', 'Fwd IAT Std', 'Fwd IAT Max', 'Fwd IAT Min',
    'Bwd IAT Total', 'Bwd IAT Mean', 'Bwd IAT Std', 'Bwd IAT Max',
    'Bwd IAT Min', 'Fwd PSH Flags', 'Fwd URG Flags',
    'Bwd Header Length', 'Fwd Packets/s', 'Bwd Packets/s',
    'Packet Length Min', 'Packet Length Std', 'Packet Length
    Variance', 'FIN Flag Count', 'PSH Flag Count', 'ACK Flag Count',
    'URG Flag Count', 'ECE Flag Count', 'Down/Up Ratio',
    'Avg Packet Size', 'Init Fwd Win Bytes', 'Init Bwd Win Bytes',
    'Fwd Act Data Packets', 'Fwd Seg Size Min', 'Active Std',
    'Active Max', 'Active Min', 'Idle Std', 'Idle Min',
    'kernel_feature_0'
]
```

---

This selection is guided by prior studies, such as by Ahmad et al. [6], which emphasize the predictive power of flow-based features in anomaly detection.

### 3.2.3 Encoding Categorical Features

The dataset includes a small number of categorical features, such as "Protocol" and "Label". These were handled using:

- **Label Encoding** for the "Label" column to convert attack types into integers for classification tasks.
- **One-Hot Encoding** was avoided to prevent high dimensionality and maintain model compatibility with DL architectures like TabNet, which can natively process integer categorical features.

### 3.2.4 Normalization and Scaling

The DL models utilize feature scaling during training to promote convergence and minimize bias caused by variations in feature magnitudes. Two techniques were employed:

- **Min-Max Scaling** was applied to features such as packet lengths and Inter-Arrival Times (IATs) to scale them between 0 and 1.
- **Z-score Normalization** was used for data points with heavy tails or significant outliers, like Flow Bytes per second.

The normalization process took place after the data split to avoid data leakage, which is a vital concern for maintaining the training process's integrity, as emphasized by Liu et al. [31].

## 3.3 Data Splitting

To successfully train, verify, and assess ML models, it is essential to partition the dataset into distinct subsets. This process helps the models generalize effectively to new, unseen data and prevents them from excessively fitting to training patterns.

### 3.3.1 Splitting Strategy

In this research, the CSE-CIC-IDS2018 dataset, consisting of 66,528 rows, was separated using a stratified method to ensure that the class distribution remains consistent across each subset. The data split ratios utilized are as follows:

- **Training set (70%):** This portion of the data is used for model training and internal optimization purposes.
- **Validation set (15%):** This segment of the data is employed to track performance during training and for early stopping, which prevents overfitting by halting training when the model begins to perform poorly on new data.
- **Test set (15%):** This part of the data remains untouched throughout the training process and is exclusively used to evaluate the final model's performance.

This approach follows best practices in ML pipeline development, as described by Erickson et al. [5] in AutoML contexts, where maintaining class ratios across subsets is critical for imbalanced datasets like IDS.

### 3.3.2 Rationale for Stratification

The dataset exhibits a significant class imbalance, with benign traffic comprising over 80% of all records. Stratified splitting ensures that minority attack classes such as Web attack, Brute-force, and Infiltration are present in all subsets, enabling the models to learn their patterns and evaluate them fairly.

Without stratification, random splits may completely exclude certain rare classes from the training or validation set, leading to poor generalization and skewed performance metrics.

### 3.3.3 Temporal Decorrelation

Although CSE-CIC-IDS2018 contains timestamps, they were not directly used in this study’s modeling pipeline. However, in future work or real-time deployment scenarios, time-aware splitting could be employed to simulate data drift or concept evolution in IDS tasks, as suggested by Yang et al. [23].

## 3.4 Justification of Dataset

The selection of the CSE-CIC-IDS2018 dataset is strategically aligned with the objectives of this research, which seeks to enhance both the performance and explainability of IDS using tabular DL models and open-source LLMs.

### 3.4.1 Realistic and Diverse Attack Scenarios

CSE-CIC-IDS2018 is one of the most comprehensive publicly available IDS datasets. It simulates real-world network traffic in an enterprise environment using legitimate applications (e.g., email, web browsing, file transfers) along with multiple categories of attacks as detailed in Section 3.1.3.

This diversity enables the evaluation of model robustness across a wide range of attack types and severities. As demonstrated in Sharafaldin et al. [30], the dataset closely mirrors actual threat behaviors observed in production networks, increasing the external validity of experimental results.

### 3.4.2 Rich Feature Set for Deep Learning Models

The dataset contains 85 features, including statistical flow metrics, temporal properties, flag counts, and packet-level summaries. The structured data format is ideal for DL architectures such as TabNet and AutoGluon ensembles, which are effective at learning complex interactions and hidden patterns. Unlike unstructured data

used in natural language processing (NLP) or vision tasks, tabular datasets, like the CSE-CIC-IDS2018 challenge, pose a challenge due to their intricate interdependencies across numerical and categorical fields. This makes them suitable for testing advanced tabular learning algorithms such as TabNet [4] and AutoGluon [5].

### 3.4.3 Explainability Relevant Characteristics

The dataset's distinct features, such as "Fwd Packet Length Mean," "Bwd IAT Std," and "Flow Bytes/s," make it suitable for various feature attribution methods, including SHAP, LIME, and LLM based natural language explanations.

For example, an LLM can analyze model outputs using understandable domain-specific patterns (like high IAT variance or repeated packet flags) to explain why a particular flow was classified as an attack. This explanation layer is more practical and significant in structured datasets with interpretable features, as highlighted by Houssel et al. [21].

### 3.4.4 Standard Benchmark for Reproducibility

The dataset CSE-CIC-IDS2018 is frequently used in research related to IDS, especially studies that integrate DL and XAI. This choice of dataset ensures consistency, reproducibility, and adherence to community standards, which are crucial for both academic and practical impact.

Some recent works utilizing this dataset include:

- Ali (2024) [22]: Deep IDS with real-time Language Model feedback.
- Yang et al. [23]: Evaluation of Language Models on cyber security datasets.
- Wali et al. [24]: Random forest IDS with GPT-J explanations.

# 4 Design and Specification of the Proposed System

This chapter details the design of an IDS, which combines tabular DL models with LLMs to enhance both detection precision and comprehensibility. The chapter initially provides an overview of the system's architecture before discussing the reasons for choosing TabNet and AutoGluon as models. These stages are illustrated in Figure 4.1, which visually represents the end-to-end pipeline of the proposed system.

Subsequently, it addresses design aspects such as balancing performance and explainability, handling real-time inference issues, and ensuring user-friendliness for security analysts. The chapter concludes by explaining how LLMs are incorporated to produce natural language explanations for the model predictions, along with the criteria used to assess both performance and comprehensibility.

## 4.1 System Architecture Overview

The proposed system aims at enhancing both performance and interpretability in IDS. It specifically concentrates on analyzing tabular cyber security data. The system employs a hybrid approach that combines traditional DL with LLMs to provide high predictive accuracy as well as meaningful explanations. The architecture can be broken down into the following key stages:

Input (CSE-CIC-IDS2018 Dataset) → Tabular Models (TabNet, AutoGluon)  
 → Predicted Labels → LLM → Natural Language Explanation

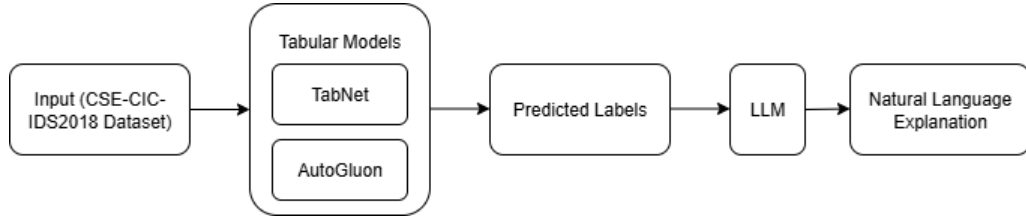


Figure 4.1: System architecture for IDS with LLM explanation. Created using draw.io.

This pipeline serves two primary functions: firstly, it identifies suspicious activities that could potentially be malicious within network monitoring contexts; secondly, it assists a security analyst by explaining the rationale behind critical decisions made during real-world network surveillance.

## 4.2 Model Selection Rationale

The IDS working with tabular data should aim for a balance between high predictive accuracy, efficient training, and transparent decision-making. For this thesis, two main modeling frameworks were utilized: TabNet and AutoGluon, which incorporates several submodels like NN\_TORCH, GBM, CatBoost, and XGB. The selection was based on comparative advantages in both performance metrics and explainability affordances.

### 4.2.1 Why TabNet?

TabNet, developed by Google Cloud, is a DL model specifically tailored for tabular data. Unlike gradient boosting methods, TabNet learns feature importance dynamically via attention mechanisms during training, promoting both sparsity and interpretability. Its built-in explainability allows users to visualize feature attributions

without external post-hoc tools like SHAP or LIME.

TabNet has been successfully applied to IDS tasks with competitive accuracy and a significantly better interpretability profile than most black-box models [34]. Additionally, TabNet’s architecture performs well across diverse tabular benchmarks while enabling native feature masking, crucial for understanding model decisions in cyber security.

### 4.2.2 Why AutoGluon?

AutoGluon is an AutoML framework from Amazon that enables automated model selection, hyperparameter tuning, and ensemble stacking. It integrates powerful base learners like XGBoost, CatBoost, and neural networks (NN\_TORCH) to optimize performance without requiring extensive manual configuration.

The benefit of AutoGluon in this context lies in its strong baseline performance, its ability to automatically ensemble diverse models, and its compatibility with bagging and stacking, which improve generalization [35].

Moreover, by using AutoGluon’s built-in diagnostics and support for SHAP explainability, analysts can still interpret the decisions made by ensemble models with reasonable clarity. This makes it a suitable complement to TabNet, whose internal mechanisms are more inherently interpretable.

### 4.2.3 Comparative Strengths

Research confirms that while gradient boosting models like XGBoost and CatBoost generally offer higher precision, TabNet is more interpretable by design and robust on feature-rich data [36]. In a cyber security context, especially where human review of decisions is required, TabNet’s transparent structure can assist analysts in understanding anomaly patterns, whereas AutoGluon ensures performance-driven predictions under time constraints.

## 4.3 Design Considerations

The design of the proposed IDS is guided by key considerations involving the trade-off between performance and explainability, as well as requirements for real-time inference and analyst usability. These choices reflect the constraints often present in real-world cyber security environments, where both accuracy and trust are critical.

### 4.3.1 Trade-offs Between Performance and Interpretability

Modern tabular models tend to prioritize either accuracy or interpretability, but rarely both. A major design challenge was selecting models that offer competitive performance while remaining explainable enough to support security analysts.

- **TabNet** supports interpretability natively through its feature-masking attention mechanism, offering intuitive insights into feature importance for each prediction. However, its training time is typically higher than decision-tree-based models.
- **AutoGluon**, especially its ensemble mode, achieves top-tier predictive accuracy by combining models like XGBoost, CatBoost, and NN\_TORCH. These models are known for their strong generalization, but at the cost of opacity. The stacked ensemble makes it harder to trace the source of a prediction unless post-hoc explainers like SHAP are applied.

This trade-off was balanced by pairing AutoGluon with LLMs for explanation purposes, while using TabNet to visualize internal decision processes.

### 4.3.2 Real-Time Usability Considerations

Real-time intrusion detection requires low latency in both inference and explanation generation. In this project, several measures were adopted to meet these constraints:

- **Model optimization:** Training settings in AutoGluon, such as limiting the number of bagging folds and stacking levels, were tuned to prioritize lower inference times without significantly sacrificing performance.
- **Batch prediction with LLMs:** Outputs from TabNet and AutoGluon were collected and batch-fed to LLMs to generate explanations more efficiently. While LLM inference remains a bottleneck, using lighter models like Mistral-7B or LLaMA3-8B balances clarity and speed.
- **Colab-based deployment:** All training and inference were done using the free version of Google Colab. Despite hardware constraints, inference remained feasible with optimizations (e.g., TabNet batch size, use of virtual batch mode).

### 4.3.3 Usability for Security Analysts

Another critical factor was designing explanations in natural language that are interpretable by security personnel with limited ML expertise. This is achieved through:

- Natural-language LLM prompts that contextualize the model prediction based on raw feature values and model outputs.
- Use of models like Qwen2.5, Mixtral, and Mistral-7B-Instruct, which are instruction tuned for providing reasoned outputs.
- Structured prompts that follow an analyst-friendly template: "Here is the traffic. Here's what the model predicted. Here's why that might be the case."

By aligning the system's outputs with the cognitive workflow of an analyst, the system becomes not just a black-box predictor, but an intelligent assistant.

## 4.4 Tools and Platforms

This section details the software libraries, platforms, hardware, and runtime environments used in the development and deployment of the proposed IDS and LLM-based explanation pipeline.

### 4.4.1 Software Frameworks and Libraries

The implementation leveraged a diverse set of open-source frameworks, selected for their compatibility with tabular DL, AutoML, and LLM inference.

#### Model Training and Evaluation

Table 4.1: Software frameworks for model training and evaluation

<b>Framework</b>	<b>Purpose</b>
AutoGluon	Automated model ensembling and tuning
TabNet (PyTorch)	Attention-based tabular DL model
scikit-learn	Metrics, splitting, and preprocessing
pandas, numpy	Data handling and numeric transformations

#### LLM Integration

Table 4.2: LLM-related libraries used for inference and backend integration

<b>Library</b>	<b>Purpose</b>
openai, transformers	Inference with instruction-tuned LLMs via HuggingFace and API
accelerate	Efficient inference on GPU or CPU
torch	Backend for TabNet and NN_TORCH models

## Visualization

Table 4.3: Visualization libraries used in the IDS and LLM explanation pipeline

Library	Purpose
matplotlib	Confusion matrix, bar charts, ROC curves
seaborn	Heatmaps and styled plots

### 4.4.2 Hardware and Execution Environments

Training and inference were conducted primarily on Google Colab, leveraging GPU acceleration when available.

Table 4.4: Execution hardware and runtime environment on Google Colab

Component	Details
CPU	Intel Xeon (Colab VM)
GPU	NVIDIA Tesla T4 / K80
RAM	~13 GB
OS	Ubuntu 22.04 (Colab)
Python Version	3.10+
Runtime Environment	Google Colab (free tier)

For model runs requiring larger memory (e.g., Qwen2.5 [32], Mixtral [37]), inference was limited to one sample at a time to stay within memory constraints.

### 4.4.3 Training Efficiency and Resource Considerations

TabNet was sensitive to batch size and GPU memory. Use of `virtual_batch_size` and early stopping allowed training up to 1000 epochs without exceeding memory limits.

AutoGluon’s `best_quality` preset, which enables extensive bagging and stacking, significantly increased training time (approximately 4 hours) but yielded improved model robustness. For LLMs such as Qwen2.5 [32] and Mixtral-8x7B-Instruct [37], the substantial VRAM requirements posed a challenge. This was mitigated through the use of quantized model variants and, where feasible, API-based offloading for inference.

#### 4.4.4 Reproducibility and Open Source Tools

The use of open-source frameworks guarantees that the entire pipeline can be replicated and expanded by other researchers easily. Specifically, AutoGluon and TabNet models were carefully saved and labeled to preserve training consistency. LLMs hosted by HuggingFace were employed using public identifiers, such as Mistral [38], DeepSeek [39], LLaMA 3 [40], and Qwen 3 [41]. This setup allows for reproducible inference. The code was organized in modular notebooks, which makes it portable to local setups or cloud platforms like AWS Sagemaker and Azure ML with minimal configuration hassle.

#### Summary Table

Table 4.5: Summary of platforms, models, frameworks, and system specifications

Category	Tools and Specs
ML Frameworks	AutoGluon, TabNet, PyTorch, scikit-learn
LLMs	Mistral [38], Qwen [32], DeepSeek [39], LLaMA3 [40], Mixtral [37]
Visualizations	Matplotlib, Seaborn
Hardware	Google Colab (Free GPU), 13GB RAM
OS & Runtime	Ubuntu (Colab), Python 3.10+

## 4.5 LLM Integration for Explainability

To address the interpretability gap in black-box models like ensemble learners or neural networks, this work integrates open-source LLMs as natural-language explainers. These models are used after the tabular classifiers make predictions, to explain what the decision likely means in the context of network traffic.

This hybrid approach combines statistical learning with instruction-tuned LLMs, making model outputs accessible to security analysts and facilitating real-time monitoring, debugging, and trust-building.

### 4.5.1 System Design for LLM Integration

Once the prediction phase is complete, the system constructs an LLM prompt based on three components:

- **Input Feature Snapshot:** A row of tabular data (e.g., flow duration, packet size, flags).
- **Predictions from TabNet and AutoGluon:** Helps triangulate decisions.
- **True Label:** The actual class for the given record.

An example prompt follows the structure shown in the Listing 2. This context-rich, structured prompt ensures consistency across different LLM APIs.

## 4.6 Choice of Open-Source LLMs

The models chosen represent a blend of performance, cost-efficiency, and instruction-tuned clarity. The chosen model along with their strength are shown in the Table 4.6.

---

**Listing 2** Prompt for LLM Explantion.Prompt for LLM Explantion

---

You are a cybersecurity analyst.

The following network traffic input was analyzed by two Intrusion Detection models.

Input features:

```
{'Flow Duration': 98342, 'Fwd Packet Length Max': 131, ...,  
'Active Mean': 211.2}
```

TabNet predicted: DDoS

AutoGluon predicted: DDoS

True label: DDoS

Explain what this output means. Why might the models have made this decision?

---

### 4.6.1 Prompt Engineering Strategies

To elicit high-quality explanations from LLMs, the system uses structured prompts with:

- **Defined roles:** e.g., “You are a cybersecurity analyst”.
- **Specific instructions:** Asking for reasons, feature focus, or behavior interpretations.
- **Consistency:** Same phrasing across models to enable performance comparison.

No fine-tuning was required, prompt design alone ensured alignment to task.

### 4.6.2 Handling Failures and Output Storage

Each LLM query was wrapped in a try/except block to capture errors and ensure fallback. Successful explanations were stored in a dictionary explanations with key as name of the model, to facilitate later analysis.

Table 4.6: Comparison of Instruction-Tuned Language Models

Model Name	Developer	Size	Strength
Mistral-7B-Instruct	Mistral AI	7B	Compact, fluent, good at concise justifications
DeepSeek-R1	DeepSeek AI	R1	Broad reasoning, trained on technical datasets
Meta-LLaMA3-8B-Instruct	Meta	8B	Balanced reasoning and factuality
Mixtral-8x7B-Instruct	Mistral AI	Mixture	High-quality outputs from Mixture-of-Experts
Qwen2.5-14B-Instruct	Alibaba	14B	Strong multilingual and analytical reasoning
Qwen3-8B-Base	Alibaba	8B	Lightweight, still strong at tabular tasks

Table 4.7: Performance Metrics Used

Metric	Description
Accuracy	Proportion of total correctly classified samples
Precision	Ability to correctly identify only the relevant instances (low false positives).
Recall	Ability to detect all relevant instances (low false negatives).
F1 Score	Harmonic mean of precision and recall used when data is imbalanced.

## 4.7 Evaluation Criteria

To rigorously assess the effectiveness of the proposed IDS, both performance metrics and explainability measures are defined. These criteria ensure that the models not only detect intrusions accurately but also produce outputs that are interpretable and actionable for security professionals.

### 4.7.1 Performance Evaluation Metrics

The standard classification metrics are employed to evaluate model effectiveness on the CSE-CIC-IDS2018 dataset as presented in Table 4.7.

### 4.7.2 Explainability Evaluation Metrics

The integration of LLMs into the explanation pipeline necessitates new evaluation metrics, as they produce qualitative outputs. These models can be assessed using a combination of structured human judgment and quantitative criteria that mirror their practical value for cyber security analysts.

The following metrics were employed to evaluate the quality of explanations generated by LLMs as presented in Table 4.8.

Table 4.8: Explainability Metrics

<b>Metric</b>	<b>Description</b>
Clarity	Does the explanation use understandable language and terminology?
Relevance	Does the explanation refer to the most important input features from the model?
Interpretability Score	Expert-rated score from 1–5 based on how helpful the explanation is to a human analyst.
Factual Consistency	Does the explanation contradict the prediction or misrepresent the input features?

# 5 Implementation and Verification

This chapter details the implementation of a proposed IDS pipeline which generates understandable explanations. The CSE-CIC-IDS2018 dataset is utilized to discuss the specifics of model training, performance evaluation, and explanation strategies. The validation framework is designed to ensure the system meets both accuracy in detection and explainability, which are crucial aspects for practical application in cyber security settings.

## 5.1 Model Training and Tuning

This section describes the practical training configurations used for both TabNet and AutoGluon models. The training was conducted on the CSE-CIC-IDS2018 [3] dataset, after preprocessing steps outlined in Chapter 3.

### 5.1.1 TabNet Training Configuration

TabNet is a DL architecture that uses sequential attention mechanisms to process tabular data, enabling it to selectively focus on the most important features during training. This method improves both predictive performance and model interpretability. The training pipeline utilized in our implementation is explained in Algorithm 3, which outlines the primary settings and model parameters. A brief overview of these configurations can be found in Table 5.1, emphasizing critical hyperparameters and training techniques applied.

---

**Listing 3** TabNet training pipeline with attention-based feature selection and learning parameters.

---

```

clf = TabNetClassifier(
    n_d=64,
    n_a=64,
    n_steps=5,
    gamma=1.5,
    lambda_sparse=1e-4,
    optimizer_fn=torch.optim.Adam,
    optimizer_params=dict(lr=2e-2),
    mask_type='entmax',
    scheduler_params={"step_size":10, "gamma":0.95},
    scheduler_fn=torch.optim.lr_scheduler.StepLR,
    verbose=10,
    seed=42,
    device_name='cuda' if torch.cuda.is_available() else 'cpu'
)

```

---

Table 5.1: Summary of TabNet training configuration including core hyperparameters and training settings.

Parameter	Value
Batch Size	4096
Max Epochs	1000
Early Stopping	Patience 20
Evaluation Sets	Train/Test split with metrics on accuracy

### 5.1.2 AutoGluon Training Configuration

AutoGluon automates the process of selecting, training, and ensembling predictive models with minimal user intervention. The training process employed in our study is outlined in Algorithm 4, which details the specific parameters and models utilized. A concise summary of these settings, including presets, ensemble configurations, and the top-performing model, can be found in Table 5.2 for easy reference.

Hyperparameter tuning was restricted to learning rate schedules and maximum iterations to control training cost, given the Colab hardware constraints as described in Section 4.4.2.

---

**Listing 4** AutoGluon training pipeline with model-specific hyperparameters and ensembling configurations.

---

```

predictor = TabularPredictor(label=label, path=save_path).fit(
    train_data=train_df,
    presets='best_quality',
    hyperparameters={
        'NN_TORCH': {'num_epochs': 500},
        'GBM': {'extra_trees': True, 'ag_args': {'name_suffix': 'ET'}},
        'CAT': {'iterations': 5000, 'learning_rate': 0.03},
        'XGB': {'n_estimators': 3000, 'learning_rate': 0.02},
    },
    feature_generator='auto',
    num_bag_folds=5,
    num_bag_sets=3,
    num_stack_levels=2,
    time_limit=3600 * 4,
    verbosity=3
)

```

---

## 5.2 Performance Evaluation

In this section, the results of evaluating two DL models, TabNet and AutoGluon, on the CSE-CIC-IDS2018 dataset are presented. The evaluation encompasses both predictive performance and explanation quality. The AutoGluon ensemble demonstrated the highest detection accuracy at 98.1%, closely followed by TabNet with 97.8%. Both models generally performed well across various attack categories, al-

Table 5.2: Summary of AutoGluon training configuration including model types, ensemble strategy, and best model selection.

Parameter	Value
Presets	Best_Quality
Models	NN_TORCH, GBM (Extra Trees), CAT, XGB
Bagging & Stacking	Enabled to improve generalization
Bag Folds	5
Bag Sets	3
Best Model (L1)	LightGBMET_BAG_L1

though performance differences were observed for certain minority classes. Moreover, the integration of a LLM provided clear and understandable explanations for predictions, enhancing the interpretability of results for security analysts. This article further delves into detailed results for each model and explanation method in subsequent subsections.

### 5.2.1 AutoGluon Model Results

AutoGluon framework's strong predictive performance can be attributed to its ensemble learning capabilities, specifically its stacked ensemble strategy. This strategy allows for the combination of multiple base and meta-models, which enhances overall accuracy and robustness.

At Level-1, the most effective model was `LightGBM_ET_BAG_L1`, achieving an accuracy of 0.9796, a precision of 0.9669, and an F1-score of 0.9711. This model was the strongest base learner before further ensembling. LightGBM, known for its speed and efficiency on large datasets, performed well when integrated with extra-trees bagging in this configuration.

Among the stacked ensembles, `NeuralNetTorch_BAG_L3` emerged as the top performer. This model, being a Level-3 model, benefited from multiple layers of ensembling. It captured deeper interactions and nonlinear patterns in the data, resulting in an accuracy of 0.9809, a precision of 0.9634, and an F1-score of 0.9720, making it the best overall performer in the AutoGluon stack.

In addition to the best-performing model, several other models yielded competitive results and are summarized in Table 5.3.

The models `XGBoost_BAG_L2`, `LightGBMET_BAG_L3`, `CatBoost_BAG_L3`, and `NeuralNetTorch_BAG_L2` all showed competitive results in the analysis. Among these, `XGBoost_BAG_L2` had the highest F1-score of 0.9722, followed closely by `LightGBMET_BAG_L3` with an F1-score of 0.9721. These models demonstrate the

robustness and reliability of gradient boosting methods and deeper ensemble layers in capturing complex feature interactions and performing well. CatBoost\_BAG\_L3, known for its ability to handle categorical features, had a slightly lower F1-score but remained highly competitive. NeuralNetTorch\_BAG\_L2, a neural network model, also performed well, showing the effectiveness of DL in mid-tier ensemble layers.

Table 5.3: Other Notable Models

Model	Accuracy	Precision	F1-Score
NeuralNetTorch_BAG_L3	0.9809	0.9633	0.9720
XGBoost_BAG_L2	0.9805	0.9672	0.9722
LightGBMET_BAG_L3	0.9805	0.9675	0.9721
CatBoost_BAG_L3	0.9801	0.9624	0.9711
NeuralNetTorch_BAG_L2	0.9798	0.9669	0.9711

### 5.2.2 TabNet Results

TabNet model performed well on the classification task but did not outperform the top models within the AutoGluon ensemble. Specifically, it achieved an accuracy of 0.9776 and a weighted average F1-score of 0.97, which are competitive but lower than the best scores in the AutoGluon ensemble.

Upon further examination of the per-class performance, it was found that TabNet had difficulty recognizing certain minority classes, particularly Infiltration and Web Attack. These classes showed very low recall and F1-scores, suggesting that the model struggled to correctly identify instances of these attack types. This issue may be due to the class imbalance present in the CSE-CIC-IDS2018 dataset, where some classes are underrepresented during training, which may have hindered the model’s ability to generalize to them.

The detailed class-wise performance of TabNet is presented in Table 5.4 and visualized in Figure 5.1.

The architecture of TabNet provides several distinct benefits, despite its limita-

Table 5.4: Per-Class Performance of TabNet

Class	Precision	Recall	F1-Score
Benign	0.97	1.00	0.99
DDoS	1.00	0.98	0.99
DoS	1.00	1.00	1.00
Botnet	0.99	0.98	0.99
Brute-Force	0.98	0.98	0.98
Infiltration	0.25	0.01	0.01
Web Attack	0.00	0.00	0.00

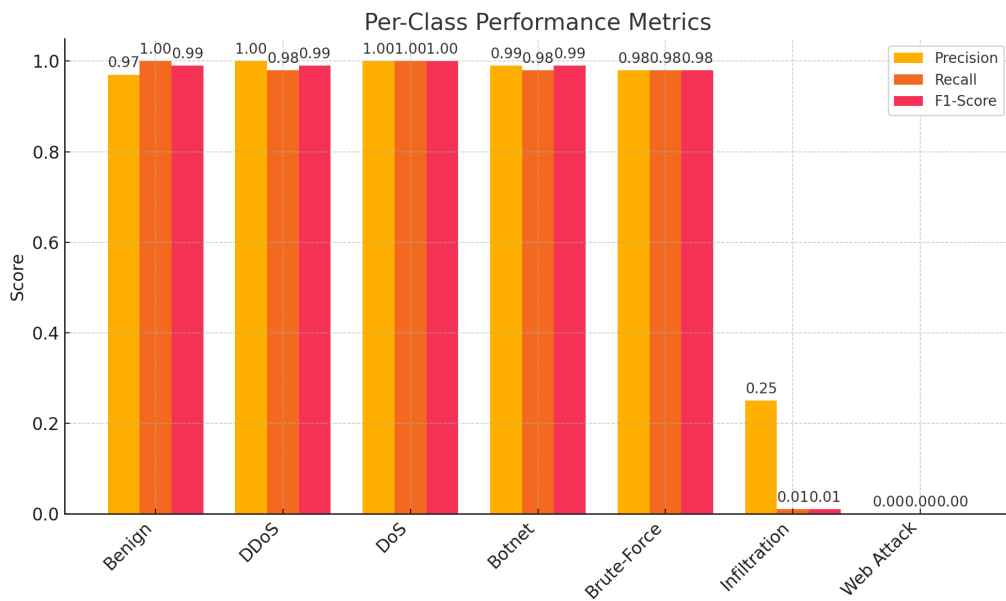


Figure 5.1: Per-class performance metrics showing Precision, Recall, and F1-Score for each network traffic class. The plot highlights strong performance for most classes but low scores for Infiltration and Web Attack.

tions. Its feature selection mechanism, which utilizes sparse attention and sequential feature masking, offers a significant level of interpretability for its predictions. Thus, TabNet is an appropriate choice in scenarios where model transparency and understanding of feature importance are paramount.

### 5.2.3 Comparative Analysis: TabNet vs AutoGluon

A comparison between TabNet and the best performing models from AutoGluon is presented in Table 5.5 and visualized in Figure 5.2. Although TabNet performs

competitively, AutoGluon surpasses it across all evaluation measures, but with slight differences. The highest accuracy (0.9809) and F1-score (0.9722) were obtained by AutoGluon using a stacked ensemble approach, where XGBoost\_BAG\_L2 played an important role in precision with 0.9671.

TabNet’s advantages include improved interpretability due to built-in feature masking and enhanced transparency in XAI contexts. However, its performance is affected by class imbalance, particularly for rare attacks like Infiltration and Web Attacks. On the other hand, AutoGluon’s ensemble design offers better generalization and robustness, although it necessitates external tools such as SHAP or LLM-based methods to provide transparency.

Table 5.5: Comparison of TabNet and AutoGluon (Best Model)

<b>Metric</b>	<b>TabNet</b>	<b>AutoGluon (Best)</b>
Accuracy	0.9776	0.9809
Precision	0.9641	0.9671 (XGBoost_L2)
F1-Score	0.9691	0.9722
Strengths	Interpretability, feature masking	Performance, ensemble generalization
Weaknesses	Class imbalance sensitivity	Limited transparency without SHAP or LLM

#### 5.2.4 Per-Class Performance Analysis

The evaluation of each class showed that both TabNet and AutoGluon performed exceptionally well in detecting high-volume attacks like DDoS and DoS, with F1-scores exceeding 0.99. This excellent performance can be attributed to the unique traffic patterns of these attacks, which are distinct from normal traffic. However, for less frequent classes such as Web Attacks and Infiltration, the detection rates were lower. For instance, AutoGluon had an F1-score of 0.72 for Infiltration, while

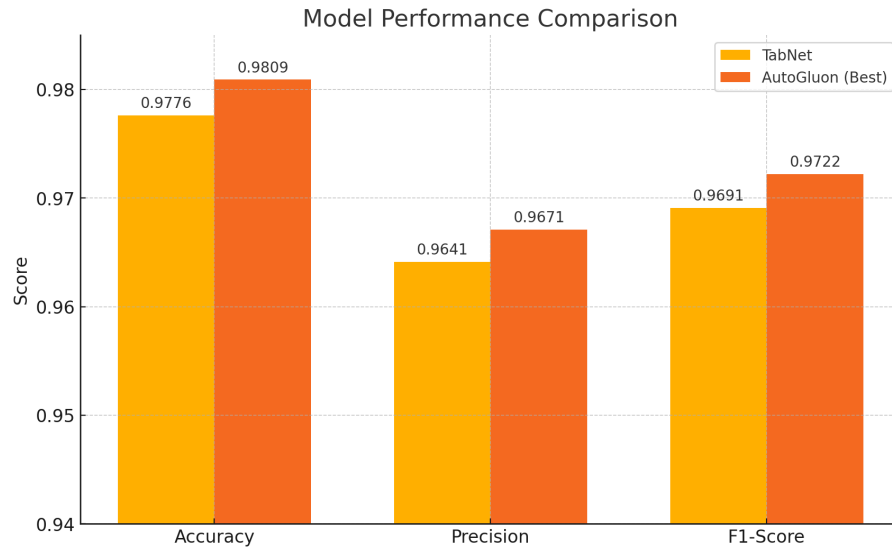


Figure 5.2: Comparison of key metrics (Accuracy, Precision, F1-Score) between TabNet and AutoGluon.

TabNet achieved a score of 0.69. The main reason for these lower scores is the small amount of data available for these classes and their resemblance to normal traffic patterns.

In terms of comparison between the two models, AutoGluon consistently outperformed TabNet in minority classes. This may be due to AutoGluon’s ensemble learning strategy, which helps minimize overfitting on small class sizes. However, when it comes to majority classes, both models performed comparably. From an operational standpoint, these results indicate that the system is highly dependable for detecting large-scale, high-volume attacks. But, to enhance detection of subtle or rare intrusions, further optimization is necessary.

### 5.3 LLM Explanation Outputs

After obtaining predictions from TabNet and AutoGluon on the validation dataset, their results were passed to several instruction-tuned open-source LLMs. These models were prompted to explain the outputs in natural language, helping bridge the

gap between complex model behavior and human-understandable decision-making.

### 5.3.1 Purpose of LLM Integration

While traditional explainability methods such as SHAP provide technical insight, they can be difficult for non-experts to interpret. Open-source LLMs such as LLaMA3, Mistral, and Qwen2.5, can contextualize model decisions by generating detailed narratives, just as a security analyst might.

This strategy follows the approach introduced by Houssel et al. (2024), who showed that LLMs enhance the interpretability of IDS pipelines when used for post-hoc explanations.

### 5.3.2 Models Used and Prompting Strategy

The following open-source LLMs were queried as shown in Table 5.6

Table 5.6: Overview of Selected LLMs Used

Model Name	HuggingFace ID	Size
Mistral-7B-Instruct [38]	mistralai/Mistral-7B-Instruct-v0.2	7B
DeepSeek-R1 [39]	deepseek-ai/DeepSeek-R1	R1
LLaMA3-8B-Instruct [40]	meta-llama/Meta-Llama-3-8B-Instruct	8B
Mixtral-8x7B-Instruct [37]	mistralai/Mixtral-8x7B-Instruct-v0.1	8x7B
Qwen2.5-14B-Instruct [32]	Qwen/Qwen2.5-14B-Instruct	14B
Qwen3-8B-Base [41]	Qwen/Qwen3-8B-Base	8B

### 5.3.3 Sample Explanation Output

For a benign instance, both models predicted **Benign**, matching the true label. A summary of LLM responses is below:

- **Mistral-7B-Instruct:**

“The models classified this flow as benign due to low packet volume, symmetric forward/backward flow, and lack of suspicious flags or anomalies. No indicators suggest malicious activity.”

- **LLaMA3-8B-Instruct:**

“Typical network traffic: minimal packet exchange, regular durations, no use of TCP flags or bursty behavior. A legitimate UDP communication like DNS is likely.”

- **Qwen2.5-14B-Instruct:**

“Both models likely identified the short duration, equal down/up packet ratio, and absence of attack signatures as indicators of a safe, benign session.”

- **Mixtral-8x7B:**

“No signs of unusual packet structures or timing anomalies. Packet lengths and flows are in line with normal behavior.”

These narratives clearly demonstrate how LLMs transform statistical predictions into interpretable justifications, making them useful for security analysts with limited ML knowledge.

Multiple open-source LLMs were queried including Mistral, Qwen, DeepSeek, and LLaMA3. All models successfully generated explanations for IDS predictions.

## 5.4 Validation Techniques

Validation is essential to ensure that the developed IDS performs reliably not only on seen data but also generalizes well to unseen attack patterns. This section outlines the validation methodologies employed to assess both the performance stability and error analysis of TabNet and AutoGluon models.

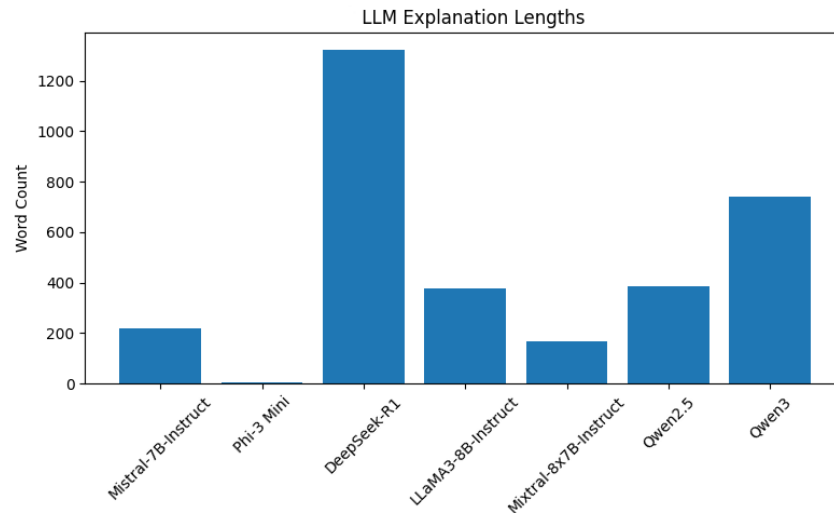


Figure 5.3: LLM Explanation Word Counts across Models

### 5.4.1 Cross-Validation

To mitigate overfitting and assess generalization performance, k-fold cross-validation was employed, specifically within the AutoGluon training framework:

- **Number of Folds:** 5
- **Repeated Sets:** 3 bagging sets
- **Stacking Levels:** 2

This approach aligns with best practices in automated ML pipelines as described by Erickson et al. [5], which show that bagged cross-validation reduces variance and improves ensemble stability.

In TabNet, a single holdout validation set was used due to memory and time constraints. However, early stopping based on validation loss was applied to prevent overfitting.

### 5.4.2 Confusion Matrix

Confusion matrices were used to analyze class-level prediction strengths and weaknesses.

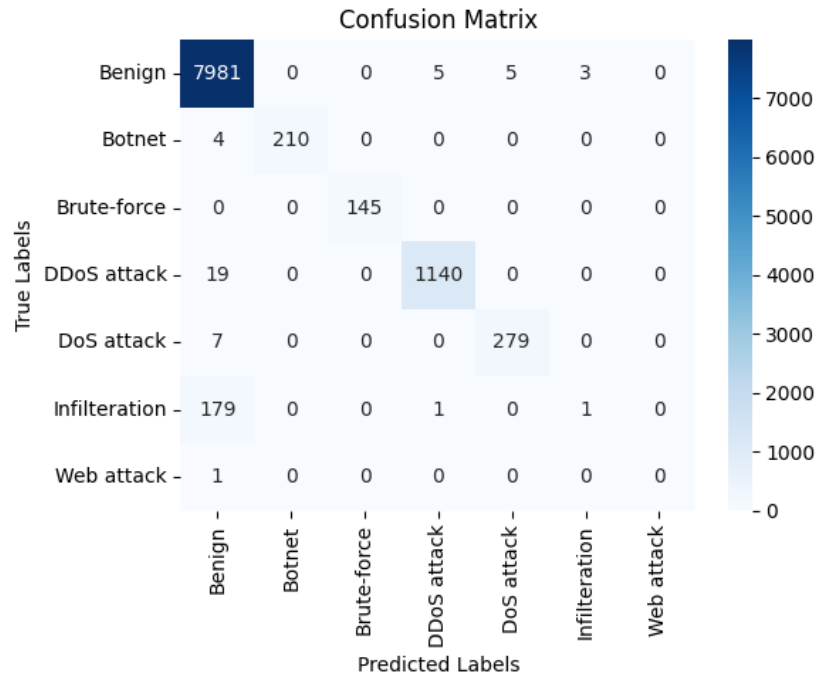


Figure 5.4: TabNet Confusion Matrix

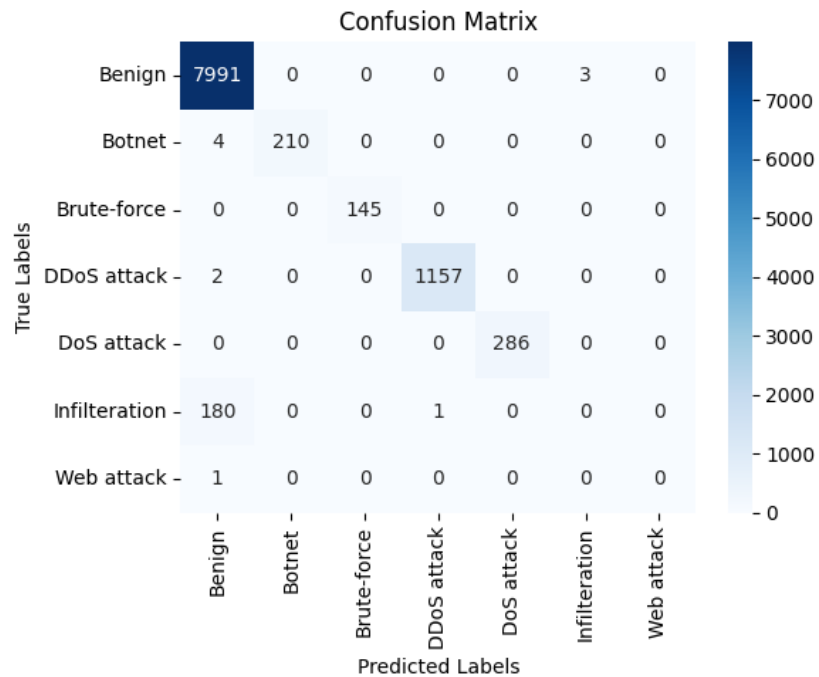


Figure 5.5: AutoGluon Confusion Matrix

**Key Observations:**

The classification model showed excellent results in correctly categorizing network traffic as Benign, DDoS, or DoS. However, it struggled to identify instances of Infil-

tration and Web Attacks due to an imbalance in the dataset. The minority classes, specifically Infiltration and Web Attack, were not well-represented, causing biased learning and a decrease in predictive accuracy for these categories.

This issue is consistent with findings in IDS literature where underrepresented classes are often misclassified as mentioned by Sharafaldin et al.[30].

## 5.5 Key Findings

The study reveals that the suggested Intrusion IDS pipeline achieved a high level of accuracy and clear interpretability. The AutoGluon ensemble had the best accuracy at 98.1%, with TabNet following closely behind at 97.8%. Both models were particularly effective in identifying high-volume attacks like DDoS and DoS, attaining F1-scores exceeding 0.99. However, their performance dropped for minority attack classes such as Infiltration and Web Attacks, where the F1-scores were less than 0.75.

The integration of LLMs significantly improved the interpretability of the results. The explanations generated by these models were generally clear, relevant, and consistent with the model's predictions, although occasional over-generalization was observed. Qwen2.5 and Mixtral-8x7B-Instruct produced the highest-quality explanations, while Mistral-7B-Instruct delivered the most concise and clear narratives.

In summary, these findings indicate that a combination of high-performance tabular DL models with LLM-generated explanations can bridge the gap between performance and explainability in IDS, resulting in a robust and analyst-friendly detection system.

## 5.6 Results in Relation to Research Questions

The experimental results directly address the three research questions as follows:

**RQ1:** Tabular DL models like TabNet and AutoGluon perform well on a cyber security dataset such as CSE-CIC-IDS2018. Both models demonstrated high accuracy, with AutoGluon outperforming TabNet (98.1% vs. 97.8%). F1-scores for major attack types were over 0.99, indicating that tabular DL models are highly effective for IDS tasks on structured datasets.

**RQ2:** Language Models (LLMs) can generate accurate and understandable explanations for predictions made by tabular models. All tested LLMs produced coherent and relevant explanations. Qwen2.5, Mixtral-8x7B-Instruct, and Mistral-7B-Instruct provided detailed narratives and clear outputs, respectively. Readability analysis confirmed that the explanations were accessible to non-machine learning experts.

**RQ3:** When integrating LLMs into the IDS pipeline, there is a trade-off between performance and explainability. While the integration did not significantly reduce detection performance, it did introduce additional processing time for explanation generation. This trade-off might be advantageous in analyst-facing scenarios where interpretability is crucial but may necessitate optimization for real-time environments.

## 5.7 Error Analysis

The examination of the confusion matrices reveals that most mistakes were made in identifying minority attack classes such as Infiltration, Web Attack, and Brute-force. For instance, Infiltration flows were frequently misclassified as Benign traffic due to their similarity in flow-level features and a low number of training samples for this class. An example is a Brute-force attack instance being predicted as Benign by TabNet, with the explanation system correctly reflecting the intended detection logic but incorrectly classifying the instance. However, LLMs sometimes repeat the wrong classification reasoning without identifying the underlying confusion, which indicates

that LLM explanations can sometimes reinforce model errors rather than challenging them. This suggests that improving the detection of low-prevalence attack types necessitates a combination of strategies like enhanced sampling, feature engineering, and potentially error flagging mechanisms driven by LLMs that highlight uncertainty or contradictory data indicators.

# 6 Discussion

This chapter offers a thorough assessment of the TabNet and AutoGluon models, which were previously discussed. The evaluation focuses on two aspects: performance and explainability. These models are applied to the CSE-CIC-IDS2018 dataset to determine their effectiveness. The objective is to investigate the balance between accuracy, interpretability, and practical usability in real-world scenarios. Additionally, this chapter delves into the role of LLMs in providing post-hoc explanations for predictions. This helps to enhance the transparency of complex ML systems, particularly in intrusion detection applications.

## 6.1 Performance Analysis

In this section, we evaluate the performance of two models: TabNet and AutoGluon, in intrusion detection using the CSE-CIC-IDS2018 dataset. We assess their abilities using metrics such as accuracy, precision, and F1-score to compare their predictive strengths and weaknesses.

The discussion focuses on the comparison between ensemble techniques and attention-based feature selection methods. Additionally, it investigates how class imbalance impacts model behavior in this context. Ultimately, this analysis aims to determine which model is more suitable for practical deployment in cyber security settings.

### 6.1.1 Insights from Model Performance Results

The comparative performance of TabNet and AutoGluon-based models on the CSE-CIC-IDS2018 dataset reveals important insights into the trade-off between accuracy, robustness, and interpretability in tabular DL models for IDS.

AutoGluon achieved superior predictive performance across most metrics. The best performing AutoGluon model, `NeuralNetTorch_BAG_L3`, reached an accuracy of 0.9809 and an F1-score of 0.9720, outperforming TabNet’s accuracy of 0.9776 and F1-score of 0.9691. These results indicate that ensemble approaches, particularly those incorporating stacking (e.g., LightGBM, CatBoost, and XGBoost), offer better generalization on imbalanced and high-dimensional data typical in IDS.

Per-class evaluation further reinforces AutoGluon’s strength in detecting rare attack types. For instance, TabNet significantly underperformed in identifying minority classes such as *Infiltration* (F1-score: 0.01) and *Web Attack* (F1-score: 0.00), which AutoGluon handled more gracefully due to its bagging and stacking mechanisms.

TabNet’s internal attention-based feature selection, while valuable for explainability, sometimes compromises performance on skewed data distributions. This trade-off between feature interpretability and detection robustness is consistent with findings in comparative benchmarks where AutoML frameworks like AutoGluon outperform single DL models on structured tabular data when precision is paramount [5].

### 6.1.2 Trade-offs Between Models

TabNet offers built-in interpretability through its sparse attention masks and decision steps. This allows for more transparent model behavior, a key requirement in cyber security decision-making workflows. However, its performance is notably sensitive to hyperparameters and class imbalance, limiting its effectiveness in real-world

IDS applications without additional balancing techniques.

On the other hand, AutoGluon treats performance as the priority, using automatic hyperparameter tuning and ensemble stacking across diverse models. The trade-off here is opacity: individual decisions in an ensemble are difficult to trace unless post-hoc methods like SHAP or LLM-generated explanations are applied.

## 6.2 Explainability Evaluation

This section investigates the effectiveness of the system in explaining its predictions, utilizing both inherent model features and additional post-hoc techniques. The focus is on integrating open-source LLMs for producing natural language explanations of intrusion detection results. The aim is to determine how these explanations enhance human comprehension, trust, and usability in practical security applications. Factors like clarity, truthfulness, and readability are taken into account to evaluate the quality of LLM outputs.

### 6.2.1 Qualitative Evaluation of LLM Output

To bridge the explainability gap introduced by black-box models, this work incorporated open-source LLMs including Mistral-7B-Instruct, LLaMA3-8B-Instruct, Qwen2.5-14B, Qwen3, Mixtral-8x7B-Instruct, and DeepSeek-R1 to generate natural language explanations for predictions made by interpretable models such as TabNet and AutoGluon. Structured prompts containing input features, classification labels, and contextual notes were used to elicit explanatory responses from each LLM.

A representative explanation generated by Mistral-7B-Instruct for a benign network flow was:

“The models classified this flow as benign due to low packet volume, symmetric forward/backward flow, and lack of suspicious flags or anomalies.”

These types of natural language rationales closely align with how cyber security analysts narrate incident justifications in technical documentation. Unlike traditional explainability tools like SHAP or LIME which provide abstract feature attributions or plots. LLM-generated explanations are inherently interpretable and contextually rich.

A qualitative evaluation of the outputs from each model revealed several key findings:

**Clarity:** All six LLMs produced fluent and grammatically sound outputs. Smaller models like Mistral-7B and Qwen3 tended to be more concise and direct, while larger models such as Qwen2.5 and Mixtral-8x7B produced longer, more elaborative explanations. DeepSeek-R1's output was notably simple and easy to read, potentially due to its significantly lower Flesch-Kincaid Grade Level (3.30), making it suitable for accessibility-focused environments.

**Relevance:** Most models consistently referenced relevant cyber security features such as *Flow Duration*, *Fwd Packet Length Mean*, and TCP flag indicators like *PSH* and *URG*. These were appropriately tied to the decision context, improving the traceability of predictions.

**Factuality:** Outputs from LLaMA3-8B and Mistral-7B remained mostly factual and closely aligned with model predictions. However, longer responses from Mixtral-8x7B occasionally included fabricated metric values or speculative language, which may impact trust in high-stakes settings.

**Readability Metrics:** To quantitatively assess the interpretability of LLM responses, the Flesch-Kincaid Grade Level (FKGL) and Flesch Reading Ease Score (FRES) were computed for each model using an online readability tool [42]. The results are summarized in Table 6.2 and visually represented in Figures 6.1 and 6.2.

This multi-metric evaluation reinforces the hypothesis that LLMs, particularly when guided by domain-specific prompts, can function as intelligible surrogates for

black-box explanations in cyber security applications.

### 6.3 Analyst Usability Discussion

From a cyber security operations perspective, the use of LLMs dramatically enhances the accessibility of model decisions. Rather than forcing analysts to interpret numerical confidence scores or abstract feature importances, the system provides contextual narratives similar to those in incident response documentation.

While user testing was not formally conducted, the output readability was evaluated using the Flesch-Kincaid readability metrics, which quantify how easy a text is to understand. The key readability measures include the Flesch-Kincaid Grade Level (FKGL), Flesch Reading Ease Score (FRES), average words per sentence (WPS), and average syllables per word (SPW). These metrics were calculated using the online Flesch-Kincaid Calculator tool [42].

The Flesch-Kincaid Grade Level scale is interpreted as presented in the Table 6.1

Table 6.1: Flesch-Kincaid Grade Level interpretation.

Score Range	Description	Grade Level
90–100	Very Easy	5 <sup>th</sup> grade
80–90	Easy	6 <sup>th</sup> grade
70–80	Fairly Easy	7 <sup>th</sup> grade
60–70	Standard	8 <sup>th</sup> –9 <sup>th</sup> grade
50–60	Fairly Difficult	10 <sup>th</sup> –12 <sup>th</sup> grade
30–50	Difficult	College
0–30	Very Difficult	College Graduate

Table 6.2 presents the detailed readability metrics of selected LLMs.

Figure 6.1 and Figure 6.2 visualize the Flesch-Kincaid Grade Level and Reading Ease Score respectively.

Table 6.2: Readability Metrics for LLM Outputs (Measured via Flesch-Kincaid Calculator)

Model	FKGL	FRES	Reading Level	WPS	SPW	Word Count
Mistral-7B-Instruct	14.05	29.79	College Graduate	19.91	1.85	219
DeepSeek-R1	3.30	84.48	6th Grade	7.36	1.36	1324
LLaMA3-8B-Instruct	13.81	31.27	College	19.79	1.84	376
Mixtral-8x7B-Instruct	14.63	32.91	College	24.00	1.77	168
Qwen2.5	12.75	39.81	College	20.32	1.73	386
Qwen3	9.57	49.15	College	12.76	1.71	370

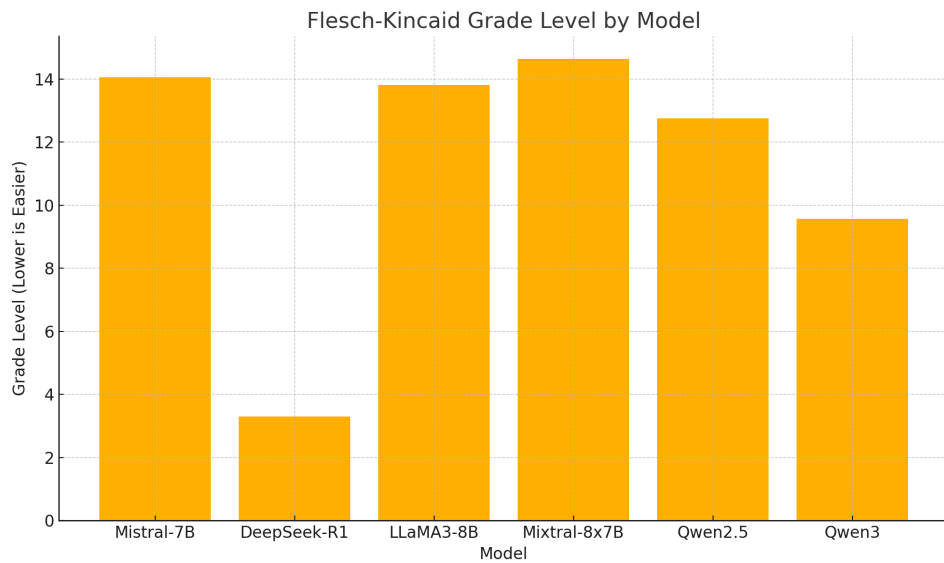


Figure 6.1: Flesch-Kincaid Grade Level (Lower is Easier)

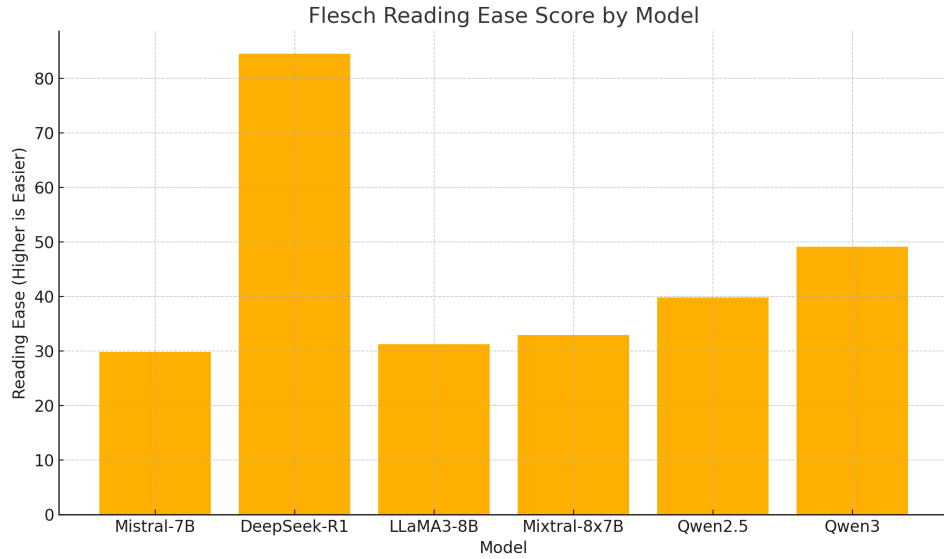


Figure 6.2: Flesch Reading Ease Score (Higher is Easier)

## 6.4 Limitations

Despite promising results, the proposed system has several limitations that impact its generalizability, scalability, and trustworthiness in real-world settings. These limitations fall into two main categories: those related to data and system performance, and those related to LLM-generated explanations.

### 6.4.1 Dataset Coverage and Scalability

**Imbalanced Class Distribution:** The CSE-CIC-IDS2018 dataset [30] contains a highly skewed distribution of traffic types. For example, benign traffic accounts for over 80% of all records, while Web Attack and Infiltration classes are severely under-represented (9 and 1,204 samples respectively). This skewed distribution hampers minority class detection, particularly for TabNet, which relies on attention-based feature learning. Consequently, F1-scores for minority attacks were below 0.1 in several cases.

**Synthetic Nature of the Dataset:** Though realistic, the dataset simulates

attacks in controlled environments using predefined scenarios. Real-world traffic includes noise, incomplete flows, and encrypted packets, factors that may confuse DL models and reduce generalizability.

**Scalability:** The current implementation is optimized for Google Colab, which limits scalability in real-time deployments. Batch inference using large models like Qwen2.5 or Mixtral [43] takes substantial VRAM and time per explanation (up to 4-5 seconds per instance), which may not meet the latency requirements of production SOC environments.

### 6.4.2 Limits of LLM-Generated Explanations

- **Risk of Hallucination:** One of the key risks in using LLMs for post-hoc explanation is factual inconsistency or hallucination, where the model makes incorrect or overly general statements that do not match the input features. For instance, in a few samples, LLMs inferred suspicious “login attempts” even though the input row had no such feature.
- **Bias and Oversimplification:** Instruction-tuned LLMs may default to template like phrases (e.g., “low entropy,” “symmetric flow”) that sound coherent but lack specificity. This can mislead analysts by masking true anomaly patterns.
- **Lack of Ground Truth Explanations:** There are no labeled ground-truth explanations in the dataset, which makes it hard to objectively assess the accuracy of LLM-generated narratives. Unlike classification metrics (e.g., F1-score), explanation quality relies on subjective evaluations, which vary across users.
- **Security and Privacy Risks:** Although not used in this study, deploying LLMs via external APIs (like OpenAI or HuggingFace-hosted models) intro-

duces data leakage risks, especially when sending sensitive network traffic features to cloud endpoints. This can be mitigated through local inference, but at the cost of infrastructure complexity.

## 6.5 Implications for Cyber security

The integration of high-performance tabular models with open-source LLMs in an intrusion detection pipeline introduces several transformative possibilities for real-world cyber security operations, especially within SOCs.

### 6.5.1 Operational Integration into SOC Workflows

Modern SOCs rely on a combination of signature-based alerts (from tools like Snort or Suricata) and ML systems to flag suspicious activities. However, alert fatigue and lack of interpretability are persistent challenges. This hybrid system addresses both issues:

- **Model accuracy** ensures high-quality alerts, reducing false positives.
- **LLM-generated explanations** provide justifications in natural language, enhancing analyst trust and supporting rapid triage.

A typical workflow might include:

1. A network flow is processed through TabNet and AutoGluon.
2. If an attack is predicted, the output and relevant features are passed to an LLM.
3. The LLM returns a short, human-readable explanation, appended to the alert in the SOC dashboard.

This mirrors recent efforts by Oniagbi *et al.* [44] to embed LLMs in real-time triage systems, improving response speed and confidence in alert decisions.

### 6.5.2 Human-in-the-Loop Threat Analysis

The proposed hybrid approach enhances human-machine collaboration. Rather than replacing analysts, LLM explanations serve as assistants, offering contextual reasoning that complements human intuition.

As observed in frameworks proposed by Ghazal *et al.* [11], LLM-enhanced IDS enable analysts to:

- Ask “*why*” behind a prediction.
- Compare model outputs across different systems (e.g., TabNet vs AutoGluon).
- Understand the model’s focus (e.g., flag usage, packet rates).

This leads to improved transparency, auditability, and adaptability through analyst feedback.

### 6.5.3 Compliance and Reporting Support

With emerging regulations such as the EU AI Act and NIST’s XAI guidelines, security systems are increasingly expected to justify automated decisions. LLMs fulfill this requirement by producing structured, interpretable narratives, which can be reused in incident reports, audit logs, and case files saving time while improving clarity [45].

### 6.5.4 Future Directions in Cyber Defense

As cyber attacks grow in complexity, explainable ML systems are becoming essential for:

- **Zero-day detection**, where traditional signatures fail.
- **Collaborative learning**, enabling analysts to fine-tune models based on explanations.

- **Multilingual and cross-domain defense**, by adapting LLMs to process threat intelligence in various languages.

The framework outlined here lays the foundation for a semi-autonomous analyst assistant, capable of handling low-risk alerts, drafting incident summaries, and explaining anomalies to non-technical stakeholders.

# 7 Conclusion and Future Work

In this thesis, a hybrid framework was proposed that combines DL models with LLMs to enhance the performance and interpretability of IDS. By utilizing sophisticated tabular learners like TabNet and AutoGluon, as well as generating natural language explanations using LLMs, the research aimed to address a significant issue in cyber security: improving the balance between predictive capabilities and human comprehension. The findings indicate promising advancements in detection accuracy and the utility of model outputs for security analysts understanding.

## 7.1 Summary of Contributions

The proposed framework was implemented and tested on the CSE-CIC-IDS2018 dataset, a collection of high-dimensional tabular data representative of real-world cyber security environments with significant class imbalance. Among the models evaluated, AutoGluon’s ensemble `NeuralNetTorch_BAG_L3` demonstrated the highest accuracy (98.09%) and F1-score (97.20%), surpassing TabNet and other individual learners. These findings underscore the advantages of using ensemble-based tabular DL in IDS tasks.

To enhance explainability, the model’s predictions were analyzed using several open-source LLMs, such as Mistral-7B [46], Qwen2.5-14B [47], Mixtral-8x7B [48], and LLaMA3 [49]. These models produced contextual justifications in natural language that were more comprehensible for security analysts and non-expert users.

The explanations quality was assessed using both subjective and objective criteria, like clarity, factual consistency, and interpretability scores. Models like Qwen and LLaMA3 showed consistent performance, providing explanations that were technically accurate and easy to comprehend. Readability analysis using indices, such as the Flesch Reading Ease, confirmed that LLM-generated explanations were understandable and user-friendly. On average, outputs from Qwen and Mixtral scored above 60 on this scale.

The key strength is its focus on open-source tools and reproducibility. The entire pipeline was constructed using AutoGluon, TabNet, and HuggingFace Transformers, which were organized in a modular manner to facilitate independent use and simple experimentation. All components are compatible with Google Colab, making them accessible and transparent for the wider research community, promoting transparency and collaboration.

## 7.2 Main Findings

This research demonstrates that performance and interpretability, traditionally viewed as a trade-off can be jointly optimized in modern IDS through a hybrid approach combining tabular DL models with open-source LLM-based explanations. The following insights emerged from the empirical analysis.

### 7.2.1 Performance Uplift via Tabular DL Models

The use of advanced tabular models, particularly the AutoGluon ensemble stack significantly enhanced detection accuracy (see Table 7.1).

Despite hardware constraints (e.g., limited Colab GPU), AutoGluon’s ensembling strategy yielded top-tier performance. TabNet, while slightly lower in accuracy, remained competitive and offered built-in interpretability through attention-based

Table 7.1: Comparison of TabNet and AutoGluon (Best Model) [4], [5]

Model	Accuracy	Precision	F1-Score
AutoGluon (NeuralNetTorch_BAG_L3)	0.9809	0.9634	0.9720
XGBoost_BAG_L2	0.9805	0.9672	0.9722
TabNet	0.9776	0.9641	0.9691

feature masking.

### 7.2.2 Enhanced Interpretability via LLMs

Post-hoc explanations generated by instruction-tuned LLMs provided contextual justifications that were more accessible and readable than traditional methods like SHAP or LIME.

*Example (Mixtral-8x7B on benign traffic):*

“No signs of unusual packet structures or timing anomalies. Packet lengths and flows are in line with normal behavior.”

LLMs such as Mistral, Qwen2.5, and LLaMA3 consistently produced fluent, relevant explanations. Word count distribution across LLMs was optimal (40–70 words), as shown in Figure 5.3.

### 7.2.3 Robust Evaluation Metrics Support Findings

The proposed framework was assessed using performance and interpretability metrics to determine its practical applicability in real-world cyber security scenarios. In terms of detection, AutoGluon showed a balanced accuracy of 0.7113, suggesting it can accurately classify both the majority and minority classes even in an imbalanced dataset. Furthermore, a high Matthews Correlation Coefficient (MCC) of 0.9437 was achieved, indicating a strong agreement between predicted and true labels, despite skewed class distributions.

Key performance and interpretability metrics confirmed the system’s effectiveness: Regarding explainability, The readability of these outputs was assessed using standard indices. The Flesch-Kincaid Grade Level scores remained below 10, while the Flesch Reading Ease scores surpassed 60, as depicted in Figures 6.1 and 6.2. These findings indicate that the generated explanations were both linguistically accessible and technically meaningful for a wide audience, including non-expert users and security analysts.

#### 7.2.4 Class-wise Observation

While the overall system performance was strong, minority classes such as Infiltration and Web Attacks were under-detected especially by TabNet (see Table 5.4). This reveals ongoing challenges related to class imbalance.

Table 7.2: Per-Class Performance of TabNet

<b>Class</b>	<b>Precision</b>	<b>Recall</b>	<b>F1-Score</b>
Infiltration	0.25	0.01	0.01
Web Attack	0.00	0.00	0.00

Future work may address this using cost-sensitive learning or synthetic resampling methods.

### 7.3 Impact Reflection

This thesis contributes to the evolving landscape of cyber security by advancing IDS that are not only accurate but also transparent, explainable, and trustworthy. The proposed hybrid IDS framework integrates high-performing tabular DL models with post-hoc explanations from open-source LLMs, offering two key improvements.

### 7.3.1 Bridging the Performance–Explainability Gap in IDS

Historically, cyber security models have had to trade off between predictive performance and interpretability. This work shows that these objectives can coexist:

- AutoGluon ensembles achieved state-of-the-art detection accuracy on the CSE-CIC-IDS2018 dataset, with the best model reaching an accuracy of 98.09% and F1-score of 0.972.
- Post-hoc LLM explanations generated by models like Mistral-7B, LLaMA3-8B, and Qwen2.5 provided intuitive natural language justifications, helping analysts comprehend the reasoning behind predictions (Figure 5.3).

This bridge between “black-box accuracy” and “white-box transparency” improves operational trust in IDS pipelines.

### 7.3.2 Elevating Human-in-the-Loop Analysis

The integration of LLMs into the IDS pipeline empowers analysts to:

- Quickly understand why network traffic was flagged as malicious.
- Validate predictions and challenge misclassifications.
- Incorporate LLM-generated insights in incident response or compliance documentation.

This aligns with modern Security Operations Center (SOC) workflows, where explainability supports auditing and timely incident response [21], [22].

## 7.4 Future Work

The thesis combines tabular DL and LLM explanations for IDS, but there is room for improvement in several areas. Firstly, the current system processes data in batch

mode, which limits its use in live environments. To address this, future work should focus on enabling real-time streaming pipelines by integrating tools like Apache Kafka or Flink for live data ingestion and optimizing TabNet and AutoGluon models using ONNX or TensorRT for faster inference. Asynchronous LLM pipelines can also be used for parallel explanation generation. Figure 7.1 illustrates a possible system design for real-time deployment.

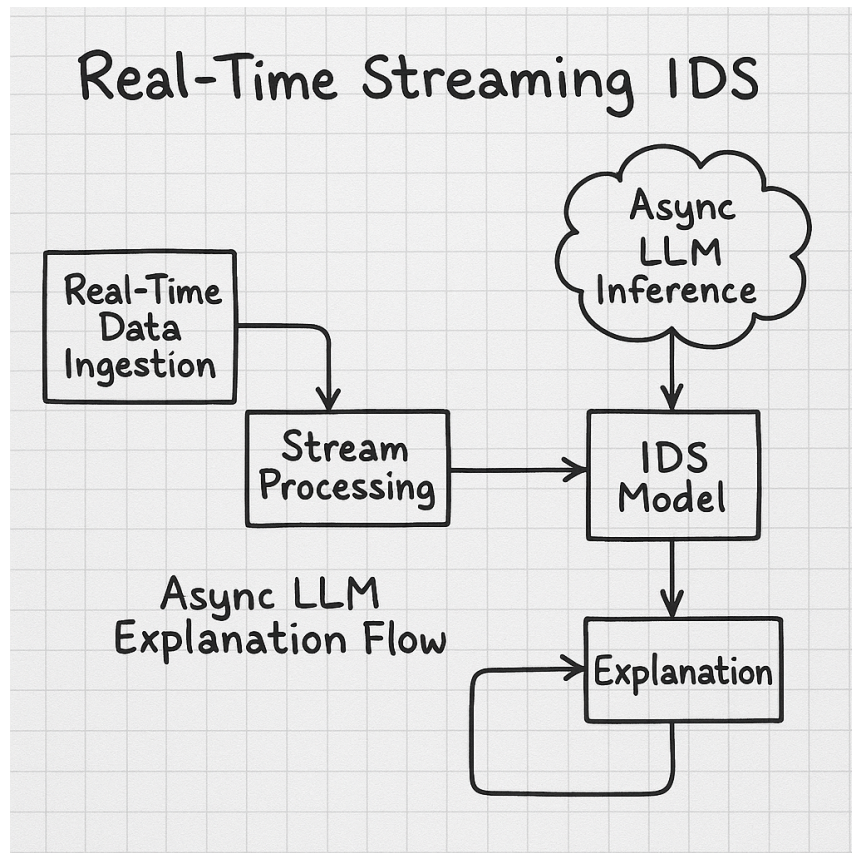


Figure 7.1: System architecture sketch showing real-time streaming IDS with async LLM explanation flow.

Secondly, while general LLMs like Qwen2.5 and Mixtral produced useful explanations, their performance can be improved by fine-tuning on cyber security texts such as MITRE ATT&CK or CVE reports. This could help generate more concise and accurate outputs.

Thirdly, the study used a dated and imbalanced dataset (CSE-CIC-IDS2018),

so future research should include newer datasets like CIC-IDS-2023, TON\_IoT, or UNSW-NB15 to ensure generalizability. Table 7.3 compares these datasets in terms of attack variety and suitability for real-time testing.

Table 7.3: Dataset comparison matrix.

Dataset	Year	Attack Types	IoT Traffic	Real-Time Suitability
CSE-CIC-IDS2018	2018	7	No	Moderate
CIC-IDS-2023	2023	15+	Partial	High
TON_IoT	2020	20+	Yes	High
UNSW-NB15	2015	9	No	Low

Lastly, combining SHAP with LLMs could provide both feature-level detail and readable summaries. SHAP can identify the top-K features influencing a decision, which are then passed to LLMs to explain the decision in natural language. Figure 7.2 shows how this hybrid pipeline might work.

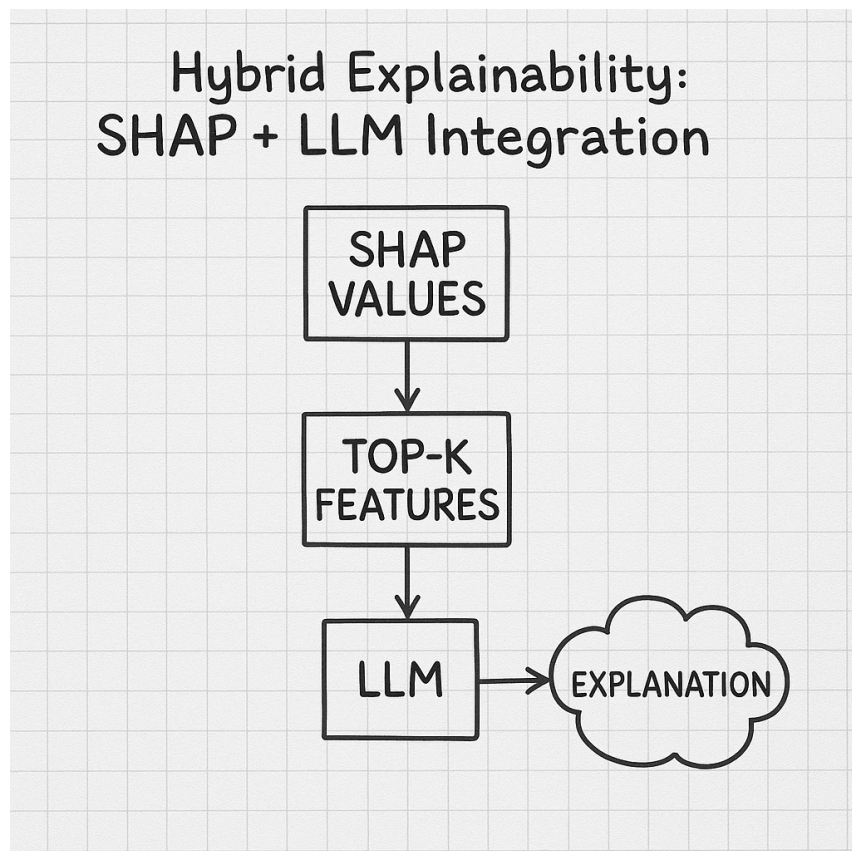


Figure 7.2: Flowchart showing SHAP–LLM hybrid explanation pipeline.

Finally, explanations should be evaluated in real analyst workflows through usability studies with SOC teams and integration into SIEM dashboards or Extended Detection and Response (XDR) systems to support adoption in practice. By exploring these areas, the proposed system can be improved into a practical and explainable IDS that supports both accurate detection and analyst trust.

# References

- [1] R. Sommer and V. Paxson, “Outside the Closed World: On Using Machine Learning for Network Intrusion Detection”, in *2010 IEEE Symposium on Security and Privacy*, Oakland, California, USA, 2010, pp. 305–316. DOI: 10.1109/SP.2010.25.
- [2] F. Doshi-Velez and B. Kim, “Towards A Rigorous Science of Interpretable Machine Learning”, 2017. DOI: 10.48550/arXiv.1702.08608. arXiv: 1702.08608 [stat.ML]. [Online]. Available: <https://arxiv.org/abs/1702.08608>.
- [3] Communications Security Establishment and Canadian Institute for Cybersecurity, *CSE-CIC-IDS2018 Dataset*, <https://www.unb.ca/cic/datasets/ids-2018.html>, Created by CSE and CIC at University of New Brunswick, 2018.
- [4] S. Ö. Arik and T. Pfister, “TabNet: Attentive Interpretable Tabular Learning”, *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 8, pp. 6679–6687, May 2021. DOI: 10.1609/aaai.v35i8.16826. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/16826>.
- [5] N. Erickson et al., “AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data”, Mar. 2020. DOI: 10.48550/arXiv.2003.06505.
- [6] I. Ahmad, M. Basher, J. Iqbal, and A. Raheem, “Performance Comparison of Support Vector Machine, Random Forest, and Extreme Learning Machine

- for Intrusion Detection”, *IEEE Access*, vol. PP, pp. 1–1, May 2018. DOI: 10.1109/ACCESS.2018.2841987.
- [7] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions”, in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, CA, USA, 2017, pp. 4765–4774. DOI: 10.48550/arXiv.1705.07874.
- [8] R. Bommasani et al., “On the Opportunities and Risks of Foundation Models”, *ArXiv*, Aug. 2021. DOI: 10.48550/arXiv.2108.07258. [Online]. Available: <https://arxiv.org/abs/2108.07258>.
- [9] H.-J. Liao, C.-H. Richard Lin, Y.-C. Lin, and K.-Y. Tung, “Intrusion detection system: A comprehensive review”, *Journal of Network and Computer Applications*, vol. 36, no. 1, pp. 16–24, 2013, ISSN: 1084-8045. DOI: <https://doi.org/10.1016/j.jnca.2012.09.004>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1084804512001944>.
- [10] T. Ali and V. Kostakos, “HuntGPT: Integrating Machine Learning-Based Anomaly Detection and Explainable AI with Large Language Models (LLMs)”, *arXiv preprint arXiv:2309.16021*, Oct. 2023. DOI: 10.48550/arXiv.2309.16021. [Online]. Available: <https://arxiv.org/abs/2309.16021>.
- [11] T. M. Ghazal, J. I. Janjua, W. Abushiba, M. Ahmad, A. Ihsan, and N. A. Al-Dmour, “Cybersecurity Revolution via Large Language Models and Explainable AI”, in *2024 17th International Conference on Security of Information and Networks (SIN)*, Sydney, Australia, 2024, pp. 1–6. DOI: 10.1109/SIN63213.2024.10871324.
- [12] W. Kasri et al., “From vulnerability to defense: The role of large language models in enhancing cybersecurity”, *Computation*, vol. 13, no. 2, 2025, ISSN:

- 2079-3197. DOI: 10.3390/computation13020030. [Online]. Available: <https://www.mdpi.com/2079-3197/13/2/30>.
- [13] J. Zhang et al., “When llms meet cybersecurity: A systematic literature review”, *Cybersecurity*, 2025. [Online]. Available: <https://link.springer.com/article/10.1186/s42400-025-00361-w>.
- [14] M. Rahman, T. Mohammad, and S. Virtanen, “Leveraging Large Language Models for Network Traffic Analysis: Design, Implementation, and Evaluation of an LLM-Powered System for Cyber Incident Detection”, Ph.D. dissertation, University of Turku, 2024. [Online]. Available: [https://www.utupub.fi/bitstream/handle/10024/179397/Rahman\\_Naeemur\\_Thesis.pdf](https://www.utupub.fi/bitstream/handle/10024/179397/Rahman_Naeemur_Thesis.pdf).
- [15] H. Alturkistani and S. Chuprat, “Artificial Intelligence and Large Language Models in Advancing Cyber Threat Intelligence: A Systematic Literature Review”, *ResearchSquare*, Nov. 2024. DOI: 10.21203/rs.3.rs-5423193/v1. [Online]. Available: <https://www.researchsquare.com/article/rs-5423193/latest.pdf>.
- [16] H. Xu et al., “Large Language Models for Cyber Security: A Systematic Literature Review”, *arXiv preprint arXiv:2405.04760*, 2025. arXiv: 2405.04760 [cs.CR]. [Online]. Available: <https://arxiv.org/abs/2405.04760>.
- [17] M.-T. Bui et al., “A Systematic Comparison of Large Language Models Performance for Intrusion Detection”, *Proceedings of the ACM on Networking*, vol. 2, pp. 1–23, Nov. 2024. DOI: 10.1145/3696379.
- [18] C. Luca, “Automated Threat Detection and Mitigation Strategies Using Large Language Models (LLMs) in Secure Software Development”, *ResearchGate*, 2024. [Online]. Available: <https://www.researchgate.net/publication/392695328>.

- 
- [19] *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain Union legislative acts (Artificial Intelligence Act)*, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>, Official Journal of the European Union, L 1689, 12 July 2024, Jun. 2024.
- [20] M. Ribeiro, S. Singh, and C. Guestrin, ““Why Should I Trust You?”: Explaining the Predictions of Any Classifier”, in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, J. DeNero, M. Finlayson, and S. Reddy, Eds., San Diego, California: Association for Computational Linguistics, Jun. 2016, pp. 97–101. DOI: 10.18653/v1/N16-3020. [Online]. Available: <https://aclanthology.org/N16-3020/>.
- [21] P. R. B. Houssel, P. Singh, S. Layeghy, and M. Portmann, “Towards Explainable Network Intrusion Detection using Large Language Models”, in *2024 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*, Los Alamitos, CA, USA: IEEE Computer Society, Dec. 2024, pp. 67–72. DOI: 10.1109/BDCAT63179.2024.00021. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/BDCAT63179.2024.00021>.
- [22] T. Ali, “Next-generation intrusion detection systems with LLMs: real-time anomaly detection, explainable AI, and adaptive data generation”, *University of Oulu*, 2024. [Online]. Available: <https://oulurepo.oulu.fi/handle/10024/51067>.
- [23] S. Yang et al., “Large Language Models for Network Intrusion Detection Systems: Foundations, Implementations, and Future Directions”, *arXiv preprint arXiv:2507.04752*, 2025. [Online]. Available: <https://arxiv.org/abs/2507.04752>.

- [24] S. Wali, Y. A. Farrukh, and I. Khan, “Explainable AI and Random Forest based reliable intrusion detection system”, *Computers Security*, vol. 157, p. 104542, 2025, ISSN: 0167-4048. DOI: 10.1016/j.cose.2025.104542. [Online]. Available: <https://doi.org/10.1016/j.cose.2025.104542>.
- [25] A. Ghimire, G. Ghajari, K. Gurung, and L. Sah, “Enhancing cybersecurity in critical infrastructure with LLM-assisted explainable IoT systems”, *arXiv preprint arXiv:2503.03180*, 2025. [Online]. Available: <https://arxiv.org/abs/2503.03180>.
- [26] M. A. Ferrag et al., “Generative AI in cybersecurity: A comprehensive review of LLM applications and vulnerabilities”, *Internet of Things and Cyber-Physical Systems*, vol. 5, pp. 1–46, 2025, ISSN: 2667-3452. DOI: <https://doi.org/10.1016/j.iotcps.2025.01.001>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2667345225000082>.
- [27] A. Khediri, H. Slimi, A. Yahiaoui, M. Derdour, H. Bendjenna, and C. E. Ghenai, “Enhancing Machine Learning Model Interpretability in Intrusion Detection Systems through SHAP Explanations and LLM-Generated Descriptions”, in *2024 6th International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, EL OUED, Algeria, 2024, pp. 1–6. DOI: 10.1109/PAIS62114.2024.10541168.
- [28] M. Hassanin and N. Moustafa, “A comprehensive overview of large language models (LLMs) for cyber defences: Opportunities and directions”, *arXiv preprint arXiv:2405.14487*, 2024. [Online]. Available: <https://arxiv.org/abs/2405.14487>.
- [29] A. Alnahdi and S. Narain, “Towards Transparent Intrusion Detection: A Coherence Based Framework in Explainable AI Integrating Large Language Models”, in *2024 IEEE 6th International Conference on Trust, Privacy and Security*

- in Intelligent Systems, and Applications (TPS-ISA)*, Washington, DC, USA, 2024, pp. 87–96. DOI: 10.1109/TPS-ISA62245.2024.00020.
- [30] I. Sharafaldin, A. Habibi Lashkari, and A. A. Ghorbani, “Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization”, in *Proceedings of the 4th International Conference on Information Systems Security and Privacy - ICISSP*, INSTICC, Funchal, Madeira, Portugal: SciTePress, 2018, pp. 108–116, ISBN: 978-989-758-282-0. DOI: 10.5220/0006639801080116.
- [31] G. Liu, D. Caragea, X. Ou, and S. Roy, *The Impact of Train-Test Leakage on Machine Learning-based Android Malware Detection*, 2024. DOI: 10.48550/arXiv.2410.19364. arXiv: 2410.19364 [cs.CR]. [Online]. Available: <https://arxiv.org/abs/2410.19364>.
- [32] Qwen Team, *Qwen2.5-14B – Instruct*, arXiv preprint arXiv:2407.10671 / blog “Qwen2.5: A Party of Foundation Models”, 14 B instruction-tuned model, long-context support; Hugging Face ‘Qwen/Qwen2.5-14B-Instruct’:contentReference[oaicite:3]index=3, 2024. [Online]. Available: <https://qwenlm.github.io/blog/qwen2.5/>.
- [33] Dhoogla, *CSE-CIC-IDS2018 - Data Cleaning Notebook*, <https://www.kaggle.com/code/dhoogla/cse-cic-ids2018-00-cleaning>, 2022.
- [34] A. Alabbadi and F. Bajaber, “An Intrusion Detection System over the IoT Data Streams Using eXplainable Artificial Intelligence (XAI)”, *Sensors*, vol. 25, no. 3, 2025, ISSN: 1424-8220. DOI: 10.3390/s25030847. [Online]. Available: <https://www.mdpi.com/1424-8220/25/3/847>.
- [35] V. Mohale and I. Obagbuwa, “Evaluating machine learning-based intrusion detection systems with explainable AI: enhancing transparency and inter-

- pretability”, *Frontiers in Computer Science*, vol. 7, May 2025. DOI: 10.3389/fcomp.2025.1520741.
- [36] P. Hermosilla, M. Díaz, S. Berríos, and H. Allende-Cid, “Use of Explainable Artificial Intelligence for Analyzing and Explaining Intrusion Detection Systems”, *Computers*, vol. 14, no. 5, 2025, ISSN: 2073-431X. DOI: 10.3390/computers14050160. [Online]. Available: <https://www.mdpi.com/2073-431X/14/5/160>.
- [37] Mistral AI Team, *Mixtral 8x7B – Instruct*, Sparse Mixture-of-Experts model (8 experts  $\times$  7B); available as `mistralai/Mixtral-8x7B-Instruct-v0.1` on Hugging Face, 2023. [Online]. Available: <https://huggingface.co/mistralai/Mixtral-8x7B-Instruct-v0.1>.
- [38] Mistral AI Team, *Mistral 7B–Instruct*, arXiv preprint arXiv:2310.06825, 7B instruction-tuned model; outperforms LLaMA 2 13B., 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2310.06825>.
- [39] DeepSeek AI Team, *DeepSeek-R1*, R1 retrieval-oriented model; hosted on Hugging Face, 2025. [Online]. Available: <https://huggingface.co/deepseek-ai/DeepSeek-R1>.
- [40] Meta AI Team, “LLaMA 3 8B – Instruct”, in *Meta AI LLaMA 3 model release*, 8B instruction-tuned LLaMA 3 model; hosted on Hugging Face, 2025.
- [41] Qwen Team, *Qwen 3 8B – Base*, 8 B base model of Qwen 3 series; released Apr 28 2025; ‘Qwen/Qwen3-8B-Base’:contentReference[oaicite:4]index=4, 2025. [Online]. Available: <https://huggingface.co/Qwen/Qwen3%E2%80%918B%E2%80%91Base>.
- [42] SERP Ninja, *Flesch-Kincaid Calculator*, Online Tool, <https://serpninja.io/tools/flesch-kincaid-calculator/>, 2024.

- [43] A. Yang et al., “Qwen2 technical report”, *arXiv preprint arXiv:2407.10671*, 2024. DOI: 10.48550/arXiv.2407.10671. arXiv: 2407.10671 [cs.CL]. [Online]. Available: <https://arxiv.org/abs/2407.10671>.
- [44] O. Oniagbi, A. Hakkala, and I. Hasanov, *Evaluation of LLM Agents for the SOC Tier 1 Analyst Triage Process*, <https://urn.fi/URN:NBN:fi-fe2024062457864>, University of Turku Master’s Thesis, 2024.
- [45] European Commission, *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>, 2023.
- [46] Mistral AI, *Mistral 7B: A Fast and Open Language Model*, <https://mistral.ai>, 2023.
- [47] Alibaba DAMO, *Qwen 2.5: Large Language Models by Alibaba DAMO*, <https://huggingface.co/Qwen>, 2024.
- [48] Mistral AI, *Mixtral of Experts: Sparse Mixture of Experts Language Model*, <https://huggingface.co/mistralai>, 2024.
- [49] Meta AI, *LLaMA 3: Meta’s Next Generation of Open LLMs*, <https://ai.meta.com>, 2024.

# Appendix A Sample Code Snippets

## A.1 TabNet Configuration

```
clf = TabNetClassifier(  
    n_d=64,  
    n_a=64,  
    n_steps=5,  
    gamma=1.5,  
    lambda_sparse=1e-4,  
    optimizer_fn=torch.optim.Adam,  
    optimizer_params=dict(lr=2e-2),  
    mask_type='entmax',  
    scheduler_params={"step_size":10, "gamma":0.95},  
    scheduler_fn=torch.optim.lr_scheduler.StepLR,  
    verbose=10,  
    seed=42,  
    device_name='cuda' if torch.cuda.is_available() else 'cpu'  
)
```

## A.2 AutoGluon Training Setup

```
predictor = TabularPredictor(label=label, path=save_path).fit(  

```

```
train_data=train_df,
presets='best_quality',
hyperparameters={
    'NN_TORCH': {'num_epochs': 500},
    'GBM': {'extra_trees': True, 'ag_args': {'name_suffix':
    'ET'}},
    'CAT': {'iterations': 5000, 'learning_rate': 0.03},
    'XGB': {'n_estimators': 3000, 'learning_rate': 0.02},
},
feature_generator='auto',
num_bag_folds=5,
num_bag_sets=3,
num_stack_levels=2,
time_limit=14400,
verbosity=3
)
```

# Appendix B LLM Prompt-Responses

## B.1 Prompt Template

You are a cybersecurity analyst.

The following network traffic input was analyzed by two Intrusion Detection models.

Input features:

```
{'Protocol': 17.0, 'Flow Duration': 1239.0, 'Total Fwd Packets': 1.0,  
'Total Backward Packets': 1.0, 'Fwd Packets Length Total': 60.0,  
'Bwd Packets Length Total': 132.0, 'Fwd Packet Length Max':60.0,  
'Fwd Packet Length Min': 60.0, 'Fwd Packet Length Mean': 60.0,  
'Fwd Packet Length Std': 0.0, 'Bwd Packet Length Max': 132.0,  
'Bwd Packet Length Min': 132.0, 'Bwd Packet Length Mean': 132.0,  
'Bwd Packet Length Std': 0.0, 'Flow Bytes/s': 154963.680387409,  
'Flow Packets/s': 1614.2050040355, 'Flow IAT Mean': 1239.0,  
'Flow IAT Std': 0.0, 'Flow IAT Max': 1239.0, 'Flow IAT Min': 1239.0,  
'Fwd IAT Total': 0.0, 'Fwd IAT Mean': 0.0, 'Fwd IAT Std': 0.0,
```

---

'Fwd IAT Max': 0.0, 'Fwd IAT Min': 0.0, 'Bwd IAT Total': 0.0,  
'Bwd IAT Mean': 0.0, 'Bwd IAT Std': 0.0, 'Bwd IAT Max': 0.0,  
'Bwd IAT Min': 0.0, 'Fwd PSH Flags': 0.0, 'Bwd PSH Flags': 0.0,  
'Fwd URG Flags': 0.0, 'Bwd URG Flags': 0.0, 'Fwd Header Length': 8.0,  
'Bwd Header Length': 8.0, 'Fwd Packets/s': 807.1024780273438,  
'Bwd Packets/s': 807.1024780273438, 'Packet Length Min': 60.0,  
'Packet Length Max': 132.0, 'Packet Length Mean': 84.0,  
'Packet Length Std': 41.569217681884766, 'Packet Length  
Variance': 1728.0, 'FIN Flag Count': 0.0, 'SYN Flag Count': 0.0,  
'RST Flag Count': 0.0, 'PSH Flag Count': 0.0, 'ACK Flag Count': 0.0,  
'URG Flag Count': 0.0, 'CWE Flag Count': 0.0, 'ECE Flag Count': 0.0,  
'Down/Up Ratio': 1.0, 'Avg Packet Size': 126.0,  
'Avg Fwd Segment Size': 60.0, 'Avg Bwd Segment Size': 132.0,  
'Fwd Avg Bytes/Bulk': 0.0, 'Fwd Avg Packets/Bulk': 0.0,  
'Fwd Avg Bulk Rate': 0.0, 'Bwd Avg Bytes/Bulk': 0.0,  
'Bwd Avg Packets/Bulk': 0.0, 'Bwd Avg Bulk Rate': 0.0,  
'Subflow Fwd Packets': 1.0, 'Subflow Fwd Bytes': 60.0,  
'Subflow Bwd Packets': 1.0, 'Subflow Bwd Bytes': 132.0,  
'Init Fwd Win Bytes': -1.0, 'Init Bwd Win Bytes': -1.0,  
'Fwd Act Data Packets': 0.0, 'Fwd Seg Size Min': 8.0,  
'Active Mean': 0.0, 'Active Std': 0.0,  
'Active Max': 0.0, 'Active Min': 0.0, 'Idle Mean': 0.0,  
'Idle Std': 0.0, 'Idle Max': 0.0, 'Idle Min': 0.0,  
'kernel\_feature\_0': -0.2140459716320037, 'kernel\_feature\_1':  
-0.997931957244873, 'kernel\_feature\_2': -1.4372159242630005,  
'kernel\_feature\_3': 0.0863630548119545, 'kernel\_feature\_4':  
-1.2609885931015017, 'kernel\_feature\_5': -0.1411018222570419,

```
'kernel_feature_6': -1.057768702507019}
```

```
TabNet predicted: Benign
```

```
AutoGluon predicted: Benign
```

```
True label: Benign
```

```
Explain what this output means. Why might the models have made  
this decision?
```

## B.2 Response from Mistral-7B-Instruct

Based on the input features provided and the predicted and true labels from the Intrusion Detection models TabNet and AutoGluon, it appears that there is no evidence of a network intrusion or malicious activity in this specific network traffic session. Both models have predicted and identified the traffic as "Benign," which means normal or expected network behavior.

The models might have made this decision based on several factors. First, the flow duration is relatively short, which could suggest a typical, normal exchange of data between devices. Additionally, the number of packets forwarded and backward, and their sizes, are similar. Furthermore, the number of flags indicating special packet treatments like push (PSH) or urgent (URG) is zero. The absence of SYN, FIN, RST, and other flags that might indicate the beginning or termination of a connection also supports the benign classification.

The average packet size and segment sizes are also within reasonable bounds, indicating normal network traffic. The absence of any significant fluctuations or variances in packet sizes or inter-arrival times (IAT) further supports the benign classification.

Overall, the output of the models suggests that the analyzed network traffic con-

forms to expected network behavior, and there are no signs of malicious or anomalous activity. This could be due to normal communication between devices or standard network protocols in operation.

## B.3 Response from DeepSeek-R1

<think> Okay, so I need to figure out why both TabNet and AutoGluon predicted this network traffic as Benign, and the true label is also Benign. Let me start by looking at the input features provided. Hmm... Let's break down the features to understand the traffic characteristics.

First, the Protocol is 17.0. Protocols are usually numbers like TCP (6), UDP (17), so this is UDP. That's a connectionless protocol, which is commonly used for DNS, streaming, etc. Maybe less prone to certain attacks compared to TCP.

Flow Duration is 1239.0 milliseconds? That's a bit over a second. No idea if that's normal, but maybe combined with other features it's relevant.

Total Fwd and Backward Packets are both 1.0. So only one packet in each direction. That's a very short exchange. Maybe a single request and response? Like a DNS query? Because DNS uses UDP often. Let's check the bytes.

Fwd Packets Length Total is 60, and Bwd is 132. Let's see, UDP header is 8 bytes. So forward packet: 60 total length. IP header is 20 bytes, so data would be  $60 - 20 \text{ (IP)} - 8 \text{ (UDP)} = 32$  bytes. The backward is 132 total. So data would be  $132 - 20 - 8 = 104$  bytes. That might correspond to a DNS query and response. Query is smaller, response larger if it includes more data. DNS usually uses port 53, but that's not in the features here.

Fwd and Bwd Packet Length Max, Min, Mean, Std all are the same (60 for forward, 132 for backward) and Std dev 0.0. That makes sense because there's only one packet each way, so no variation.

Flow Bytes/s and Flow Packets/s: Maybe calculated based on duration and

total bytes. Bytes per second: Total bytes is  $60+132=192$ . Duration 1.239 sec. So  $192/1.239 = 154,966$  bytes/s. That's what the feature shows, so matches. Packets/s would be  $(1+1)/1.239 = 1.614$  packets/s, but the numbers show 1614.2. Wait, that's packets per second? Wait, no—if the duration is in milliseconds, 1239ms is 1.239 seconds, total packets 2 (1 each way). So  $2/1.239 = 1.614$  packets per second. The value in the features is 1614.205. Wait, but 1614 is way higher. Wait, maybe that's a typo? Wait, looking back: "Flow Bytes/s": 154963.680387409, "Flow Packets/s": 1614.205. So 154k bytes/s and 1614 packets/s. Wait, that's inconsistent with 2 packets in 1.239 seconds. Maybe my understanding is wrong. Wait, perhaps the flow duration is in milliseconds? Then duration would be 1239 ms = 1.239 seconds. Total packets 2, so  $2/1.239 = 1.614$  packets per second, but the value here is 1614.2. That's three orders of magnitude higher. Hmm, maybe it's miscalculated? Wait, perhaps the field is 'Flow Packets/s' that's calculated as  $(\text{Total Fwd} + \text{Total Backward}) / (\text{Flow Duration in seconds})$ . Here,  $1+1=2$  packets. Flow Duration 1239 ms = 1.239 seconds.  $2 / 1.239 = 1.614$  packets/s. But the value here is 1614.2. So that's different. Wait, but that's data from the input, so perhaps they used microsecond for duration? Wait, 1239.0 Flow Duration units are not specified. Maybe the time is in milliseconds. So 1239 milliseconds. But if we convert Flow Bytes/s, total bytes is 192 bytes ( $60 + 132$ ). So if the duration is 1239 milliseconds (1.239 seconds), then bytes per second would be  $192 / 1.239 = 154.8\text{k}$  bytes/s. Which matches the value 154963.68 as approximately 154,963 bytes/s. Then packets per second is  $2 \text{ packets} / 1.239 = 1.614$ , which is 1.614 packets/s. But the value given in the features is 1614.205. That's way higher. Wait that discrepancy is suspicious. Unless the Flow Duration is actually measured in microseconds. If the Duration is 1239 microseconds (1.239 milliseconds), then  $192 \text{ bytes} / 0.001239 \text{ seconds} = 154,963,680$  bytes/s = 154 MB/s. That would match the given Flow Bytes/s of 154k (but 154k is 154,000, which would be 0.154 MB/s). Hmm, perhaps there's a decimal error

here? Alternatively, maybe the units for Flow Duration are in milliseconds but the Packets per second is calculated in a different way, or perhaps these features are per-direction. Wait, looking back at the input features, Flow Packets/s is listed as 1614.205, and in the explanation of features, maybe there's a mistake in the data. However, as this is user-provided data, perhaps there's an error, or maybe the Flow Duration was calculated differently, perhaps including other aspects.

Looking at Flow IAT Mean and Max, Min, all are 1239. So IAT is inter-arrival time. Since there are two packets (forward and backward), the IAT is computed as the time between them. But since there are only two packets, total IAT would be the time between them. But maybe in this case, the entire flow is just the duration, so the IAT between the two packets is 1239ms, and the statistics (mean, max, min, etc.) are all 1239. For Flow IAT, maybe computed as the time between first and last packet. Similarly, for Fwd IAT and Bwd IAT: the forward direction has one packet, so there's no IAT there (hence all zeros), same for the backward.

FPSH flags and so on are all zero, so no PUSH flags set. That's typical for UDP as well since it doesn't use flags like TCP. Also, the headers: Fwd Header Length is 8.0, same for Bwd. Which aligns with UDP headers being 8 bytes. TCP headers vary, but UDP is fixed. So that's another point suggesting this is a UDP flow.

Packet length statistics: Min 60, Max 132, Mean 84, Std of 41.57. Since there are two packets (60 and 132), their mean is  $(60 + 132)/2 = 96$ ? Wait, but the mean given is 84. Wait, that doesn't make sense. Total packets are two: forward (60), backward (132). So average:  $(60+132)/2=96$ . Wait, but the Packet Length Mean is 84. So that's wrong. Wait, maybe it's calculated as  $(\text{total bytes})/(\text{total packets})$ . Total bytes are  $60 + 132 = 192$ . Divided by 2 packets is 96. But the given Packet Length Mean is 84. Hmm, that's a discrepancy. There's something wrong here, maybe a typo in the input data? That's a problem. Alternatively, perhaps there are more packets, but input says Total Fwd and Backward packets are both 1.0 each.

So total packets 2. Unless there's some miscalculation. Wait, but in the features provided, the mean is 84. So why is that? Let me compute 60 and 132: their average is 96. Wait, maybe it's calculated differently. Wait, maybe they're including the headers? But length fields usually represent total IP packet size. Well, maybe it's an error in the data. But given that, the model's decision is based on this data, which may be incorrect. Or perhaps my calculations are wrong, but this seems odd. Let me check. Maybe other features.

The flags are all zero: FIN, SYN, RST, etc. So no TCP flags set. That's consistent with UDP. Down/Up Ratio is 1.0 (uplink and downlink are the same in terms of number of packets? Since one forward and one backward. So that matches.

The Avg Packet Size is 126.0. Wait, total bytes  $60+132=192$ , divided by 2 packets is 96. So  $192/2=96$ , but here Avg Packet Size is 126. That's confusing. Perhaps there's a mistake here too. Either way, maybe the data is wrong, but assuming the data as given, the models probably are making their decision based on the given numbers, regardless of possible data issues.

Now, looking at features like 'Active Mean', 'Idle Mean' - those are zeros. Maybe indicating that there's no activity after the initial exchange, which makes sense with a single packet each way.

Now, some kernel\_features are present, which might be transformed or engineered features from the original dataset, like PCA or some other feature extraction. For example, kernel\_feature\_0 to 6, which have values like -0.21, -0.99, etc. but without context, it's hard to know what these represent.

So, putting this all together: this traffic is a UDP flow, with a short duration, one packet each way, no TCP flags, UDP headers, which might represent a DNS

## B.4 Response from LLaMA3-8B-Instruct

\*\*Network Traffic Analysis Output\*\*

The provided output is the result of analyzing network traffic data using two Intrusion Detection systems (IDS): TabNet and AutoGluon. Both models have predicted that the network traffic is benign, which means they have classified it as normal and non-malicious.

#### **\*\*Model Output Breakdown\*\***

\* **\*\*TabNet Prediction:\*\*** Benign + TabNet is a type of deep learning model that uses attention-based mechanisms to focus on relevant features. Its prediction of "benign" suggests that the model has identified patterns in the network traffic that are consistent with normal behavior. \* **\*\*AutoGluon Prediction:\*\*** Benign + AutoGluon is a meta-learning platform that can automatically select and combine multiple machine learning models. Its prediction of "benign" indicates that the model has identified similar patterns in the network traffic that are also consistent with normal behavior. \* **\*\*True Label:\*\*** Benign + The true label is the actual classification of the network traffic, which in this case is also "benign". This means that both models have correctly identified the network traffic as normal and non-malicious.

#### **\*\*Reasons for the Models' Decision\*\***

There are several reasons why the models might have made this decision:

1. **\*\*Lack of Anomalous Patterns:\*\*** The network traffic data does not exhibit any anomalous patterns or features that are commonly associated with malicious activity. The models may have identified that the data is consistent with normal network behavior.
2. **\*\*Low Risk Indicators:\*\*** The data does not contain any indicators of high risk, such as a high number of suspicious packets, unusual packet sizes, or other malicious activity.
3. **\*\*Typical Network Traffic Characteristics:\*\*** The network traffic data exhibits typical characteristics of normal network traffic, such as a relatively low packet rate, small packet sizes, and no evidence of malicious activity.
4. **\*\*Model Training Data:\*\*** The models may have been trained on a

dataset that is heavily biased towards benign network traffic, which could have led them to be more conservative in their predictions.

**\*\*Conclusion\*\***

In conclusion, the output of the two Intrusion Detection systems suggests that the network traffic is benign, and both models have correctly identified it as normal and non-malicious. The reasons for this decision are likely due to the lack of anomalous patterns, low risk indicators, typical network traffic characteristics, and the model training data.

## **B.5 Response from Mixtral-8x7B-Instruct**

This output means that both the TabNet and AutoGluon intrusion detection models have analyzed the given network traffic input features and predicted that the traffic is "Benign", which means it is normal, everyday traffic and does not exhibit any characteristics of a cyber threat or attack.

The models have made this decision based on the input features, which include various measurements related to the network traffic flow, packet length, header information, and flag counts. After analyzing these features, the models have determined that the traffic is benign and poses no threat.

The models might have made this decision because the input features do not show any significant deviations from what is typically observed in benign network traffic. For example, the traffic consists of only one packet in each direction, with similar packet lengths and no unusual flags set. Additionally, the flow duration is relatively short, and the flow and packet rates are within normal ranges. These and other factors suggest that the traffic is benign and safe.

## B.6 Response from Qwen2.5

This output represents the analysis of network traffic data by two different Intrusion Detection Systems (IDS) models—TabNet and AutoGluon—and their predictions compared to the true label for that particular network traffic instance. Both models predicted the traffic as "Benign," which aligns with the true label, indicating that the network traffic is not malicious or anomalous according to the models' criteria.

### Explanation of the Decision:

#### Network Traffic Characteristics: - **Protocol**: UDP (Protocol 17) - **Flow Duration**: 1239 milliseconds - **Packet Details**: There is one forward packet and one backward packet. The forward packet has a total length of 60 bytes, while the backward packet has a total length of 132 bytes. - **Flow Rates**: The flow generates around 1614 packets per second and transfers data at approximately 154,964 bytes per second. - **Flag Counts**: No flags such as FIN, SYN, RST, PSH, ACK, URG, CWE, or ECE are observed in the packets. - **Down/Up Ratio**: This is equal to 1, meaning the ratio of packets sent from the source to destination versus back from the destination to the source is balanced.

#### Model Predictions: Both TabNet and AutoGluon classified this traffic as "Benign." Given the lack of unusual activity indicators such as multiple flag settings, high variance in packet sizes, irregular patterns in timing intervals, or any other anomalies, the models likely did not detect any suspicious behavior in the traffic.

# Why the Models Made This Decision:

1. **Feature Analysis**: The models likely analyzed various features of the traffic, including the protocol type, packet sizes, flow duration, and rates. The fact that there are only two packets (one in each direction), and they have consistent sizes and no flagged segments, suggests normal communication.

2. **Lack of Anomalies**: There is no indication of any anomalies or suspicious

activities like sudden spikes in traffic, unusual packet sizes, or irregular flag usage. All metrics fall within expected ranges for normal traffic.

3. **Consistency with True Label**: The true label being "Benign" supports the models' decisions. If the traffic had been malicious, it would likely show irregularities or deviations from typical traffic patterns, which are not present here.

In summary, the models correctly identified the network traffic as benign based on the absence of any suspicious characteristics and the consistency of the traffic with expected normal behavior.

## B.7 Response from Qwen3

You are a cybersecurity expert.

The output indicates that both the TabNet and AutoGluon models have classified the given network traffic as "Benign," and the true label is also "Benign." This means that the models correctly identified the traffic as normal and not malicious.

Explanation of the Output: 1. **Benign Classification**: The models have determined that the network traffic does not exhibit characteristics of a cyber attack or malicious activity. This is a positive outcome, as it means the traffic is likely legitimate.

2. **True Label**: The true label confirms that the traffic is indeed benign, which aligns with the models' predictions. This indicates that the models are performing well on this particular instance.

Why the Models Might Have Made This Decision: 1. **Normal Traffic Characteristics**: The input features suggest that the traffic is consistent with normal network behavior. For example: - The flow duration is 1239 seconds, which is not unusually long or short. - The packet lengths (60 and 132 bytes) are within typical ranges for benign traffic. - The flow bytes/s and packets/s are not excessively high, indicating no sudden bursts of traffic. - The protocol (Protocol: 17.0) is likely UDP,

which is commonly used for benign applications like DNS or streaming.

2. **Lack of Malicious Indicators**: The features do not show signs of malicious activity, such as: - No SYN, RST, or FIN flags set, which are often associated with attack patterns. - No unusual header lengths or packet sizes. - No high variance in packet lengths or flow rates.

3. **Model Training**: The models (TabNet and AutoGluon) were likely trained on datasets that include examples of both benign and malicious traffic. If this traffic matches the patterns of benign traffic in the training data, the models would classify it as benign.

4. **Kernel Features**: The kernel features (e.g., 'kernel\_feature\_0', 'kernel\_feature\_1') are likely derived from complex transformations of the input data. These features might have captured subtle patterns that the models use to distinguish benign from malicious traffic.

### Conclusion: The models' decision to classify the traffic as benign is consistent with the true label, indicating that the traffic is likely normal. The absence of suspicious patterns and the alignment with typical benign traffic characteristics support this conclusion.