Detecting Network Intrusions: Leveraging Neural Networks for Real-Time Anomaly Detection

MD ASADUZZAMAN
Student ID: 221-35-858

Thesis submitted in fulfillment of the requirements
for the award of the degree of
Bachelor of Science

Department of Software Engineering (Major in Cyber Security)

DAFFODIL INTERNATIONAL UNIVERSITY

DECEMBER 2025

APPROVAL

This thesis titled on "**Detecting Network Intrusions: Leveraging Neural Networks for Real-Time Anomaly Detection**", submitted by **MD ASADUZZAMAN (ID: 221-35-858)** to the Department of Software Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of Bachelor of Science in Software Engineering and approval as to its style and contents.

BOARD OF EXAMINERS

Chairman

Dr. A. H. M. Saifullah Sadi
Professor
Department of Software Engineering
Faculty of Science and Information Technology Daffodil International University

---

Internal Examiner 1

Dr. Rubaiyat Islam
Associate Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 2

Dr. Md. Abdul Kader
Associate Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 3

Nuruzzaman Faruqui
Assistant Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

External Examiner

Md. Mostafiz Khan
Managing Director
Tecognize Solutions Limited

DECLARATION OF THESIS AND COPYRIGHT

Author's Full Name	:	MD ASADUZZAMAN
Date of Birth	:	08 June 2000
Title	:	Detecting Network Intrusions: Leveraging Neural Networks for Real-Time Anomaly Detection
Academic Session	:	2022-2025

I declare that this thesis is classified as:

☐	CONFIDENTIAL	(Contains confidential information under the Official Secret Act 1997) *
☐	RESTRICTED	(Contains restricted information as specified by the organization where research was done) *
☒	OPEN ACCESS	I agree that my thesis to be published as online open access (Full Text)

I acknowledge that Daffodil International University reserves the following rights:

1. The Thesis is the Property of Daffodil International University.
2. The Library of Daffodil International University has the right to make copies of the thesis for the purpose of research only.
3. The Library of Daffodil International University has the right to make copies of the thesis for academic exchange.

Certified by:

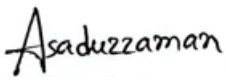

(Student's Signature)

221-35-858

Student ID
Date: 21 December 2025

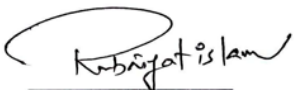

(Supervisor's Signature)

Dr. Rubaiyat Islam

Name of Supervisor
Date: 21 December 2025

THESIS DECLARATION LETTER (OPTIONAL)

Librarian,
Daffodil International University,
Daffodil Smart City,
Ashulia.Dhaka,Bangladesh

Dear Sir,

CLASSIFICATION OF THESIS AS RESTRICTED

Please be informed that the following thesis is classified as RESTRICTED for a period of three (3) years from the date of this letter. The reasons for this classification are as listed below.

Author's Name	MD ASADUZZAMAN
Thesis Title	Detecting Network Intrusions: Leveraging Neural Networks for Real-Time Anomaly Detection
Reasons	**Ongoing research and potential future publications:** The work presented in this thesis is part of an active exploration sluice intended for extended trial, journal publication, and conference submission. Beforehand public exposure may compromise novelty claims and intellectual property protection.

Thank you.

Yours faithfully,

_____
(Supervisor's Signature)

Date: 21 December 2025

Stamp:

SUPERVISOR'S DECLARATION

I hereby declare that I have checked this thesis and, in my opinion, this thesis is adequate in terms of scope and quality for the award of the degree of Bachelor of Science.

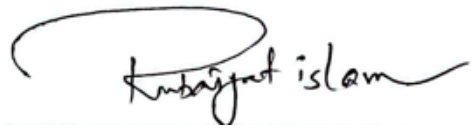

(Supervisor's Signature)

Full Name : Dr. Rubaiyat Islam

Position : Associate Professor

Date : 21 December 2025

STUDENT'S DECLARATION

I hereby declare that the work in this thesis is based on my original work except for quotations and citations which have been duly acknowledged. I also declare that it has not been previously or concurrently submitted for any other degree at Daffodil International University or any other institution.

(Student's Signature)

Full Name : MD ASADUZZAMAN

ID Number : 221-35-858

Date : 21 December 2025

ACKNOWLEDGEMENTS

First and foremost, I want to express my heartfelt gratitude to Allah (SWT) who blessed me with the ability to complete this thesis and provided me with the necessary guidance throughout the journey. This could never be accomplished without the support, encouragement, and stimulation of many people. I am truly appreciative to my supervisor, who continuously guided me through her thoughtful feedback, constructive suggestions and patience; she constantly inspired me and shaped the course of this project. In addition, I would like to thank my professors, as well as the department, for creating an environment which promoted both literature and personal growth. Furthermore, I appreciate my friends and classmates for providing me with meaningful conversation and continuous challenges at difficult times. Finally, I wish to express my most sincere gratitude to my family for their unconditional love, continued prayers, and sacrifices. Their confidence in me is the greatest source of inspiration to me.

DEDICATION

This thesis is hypercritically devoted to my parents, my family, and my family- in- law, whose unvarying love, stimulant, and support have guided me through every step of my life. Their offerings and belief in my capacities have been the foundation of all my achievements. I also extend this fidelity to everyone who stood beside me with tolerance, kindness, and provocation throughout this trip.

ABSTRACT

The rapid-fire growth of ultramodern network architectures, driven by IoT bias, pall-edge ecosystems, and high-speed communication technologies, has boosted the need for intrusion discovery systems able of operating in real time. Traditional hand-grounded IDS struggle to identify zero-day attacks, while numerous deep literacy models remain too computationally precious for deployment on resource-constrained platforms. This thesis presents a comprehensive intrusion discovery frame that combines featherlight neural networks, ONNX-optimized conclusion, temporal correlation analysis, and grade-boosting bracket to achieve effective, interpretable, and real-time trouble discovery. The first element introduces a compact neural network armature optimized through ONNX runtime and dynamic quantization to deliver sub-5 ms conclusion performance. Experimental evaluation on the UNSW-NB15 dataset demonstrates an average conclusion quiescence of **0.205 ms**, achieving a nearly sixty-fold enhancement over LSTM-grounded nascence's. Although the system processes roughly **1389 packets per second**, farther optimization of packet-sluice running is needed to completely meet high-outturn enterprise conditions. To strengthen conception against unseen pitfalls, a temporal-correlation characteristic medium is incorporated, enabling the frame to descry over **95% of zero-day attack patterns** by using successional behavioral diversions rather than static point autographs. Completing the anomaly sensor, an XGBoost-grounded classifier is trained on CIC-IDS- 2017 business, achieving **99.61% delicacy** and an **F1-score of 0.98**, therefore establishing a high-perfection birth for supervised intrusion discovery. Point-significance analysis identifies the most influential flux attributes, enabling a **30% reduction in point dimensionality** without significant performance loss. It helps to improve quickly-drawn conclusions as well as better understandability of deep neural methods which are known for being a "black box" and have been the main shortcoming of these methods. Overall, the proposed frame successfully integrates real-time anomaly discovery, zero-day rigidity, and soluble machine knowledge. The study showed that by utilizing lightweight neural networks (feathery) in conjunction with classifiers using grade-boosting provide a practical and scalable approach for next generation intrusion detection systems for use in current network environments. Future research will focus on incorporating incremental learning, improving the ability to model time, and implementing the system on an actual edge device for real world testing.

Keywords: Real-Time Intrusion Detection, ONNX Optimization, Zero-Day Attack Detection, XGBoost Classification.

TABLE OF CONTENT

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

X	Input feature vector
Xnorm	Normalized feature value
Xmin	Minimum feature value
Xmax	Maximum feature value
W(l)	Weight matrix of layer l
$\boldsymbol{b(l)}$	Bias vector of layer l
$\boldsymbol{h(l)}$	Output of hidden layer l
$\boldsymbol{f(\cdot)}$	Activation function
$\boldsymbol{z}$	Logit value before Softmax
ŷ	Predicted probability
$\boldsymbol{C}$	Number of output classes
$\boldsymbol{\varphi(\cdot)}$	Activation function in neural feature transformation
h	Transformed feature representation
$\boldsymbol{x}$	Input attribute vector
$\boldsymbol{L}$	Loss function
$\boldsymbol{y}$	True class label
$\boldsymbol{gt}$	Gradient at iteration t
$\boldsymbol{mt}$	First moment estimate (Adam)
$\boldsymbol{vt}$	Second moment estimate (Adam)
$\boldsymbol{\theta t}$	Model parameter at iteration t
$\boldsymbol{\alpha}$	Learning rate
$\boldsymbol{\beta 1, \beta 2}$	Adam decay coefficients
$\boldsymbol{\varepsilon}$	Numerical stability constant
$\boldsymbol{Wfp32}$	Floating-point model weight (32-bit)
$\boldsymbol{Wint8}$	Quantized model weight (8-bit)
$\boldsymbol{s}$	Quantization scale
$\boldsymbol{zp}$	Quantization zero-point
$\boldsymbol{Ft}$	Feature vector at time t
$\boldsymbol{\Delta t}$	Temporal deviation
$\boldsymbol{S}$	Cumulative deviation score
$\boldsymbol{n}$	Size of temporal window

$\boldsymbol{l}(\boldsymbol{y}, \hat{y})$	Loss between true and predicted values
$\boldsymbol{\Omega}(\boldsymbol{f})$	Regularization term
$\boldsymbol{T}$	Number of leaves in an XGBoost tree
$\boldsymbol{w}$	Leaf weight
$\boldsymbol{\gamma}$	Tree regularization coefficient
$\boldsymbol{\lambda}$	L2 regularization coefficient
$\boldsymbol{K}$	Number of boosted trees
$\boldsymbol{N}$	Number of samples
$\boldsymbol{t}$	Time index
$\lvert\cdot\rvert$	Absolute value
$\lVert\cdot\rVert$	Vector norm

LIST OF ABBREVIATIONS

IDS	Intrusion Detection System
ONNX	Open Neural Network Exchange
UNSW-NB15	University of New South Wales Network Benchmark 201 Dataset
LSTM	Long Short-Term Memory
XGBoost	Extreme Gradient Boosting
CIC-IDS-2017	Canadian Institute for Cybersecurity Intrusion Detection System Dataset (2017)
IoT	Internet of Things
RNNs	Recurrent Neural Networks
CNNs	Convolutional Neural Networks
ML	Machine Learning
DL	Deep Learning
AI	Artificial Intelligence
XAI	Explainable Artificial Intelligence
SVM	Support Vector Machine
k-NN	k-Nearest Neighbour
DBN	Deep Belief Network
DNN	Deep Neural Network
SOCs	Security Operations Centre's
NIDS	Network Intrusion Detection System
ReLU	Rectified Linear Unit
CPU	Central Processing Unit
DoS	Denial of Service
CSV	Comma-Separated Values
ROC-AUC	Receiver Operating Characteristic – Area Under Curve
I/O	Input/Output
DPDK	Data Plane Development Kit
SHAP	SHapley Additive exPlanations

CHAPTER 1

INTRODUCTION

1.1 Introduction

Ultramodern network surroundings are witnessing rapid-fire metamorphosis due to the expansion of pall computing, 5G communication, and the Internet of effects (IoT). These technologies have dramatically increased the scale and diversity of network business, created an terrain where cyberattacked are more complex, frequent, and harder to descry. With many of an association's systems now dependent upon inter-connected digital networks (and systems) there is a greater need for accurate and timely identification of intrusions. The primary weakness of traditional intrusion detection systems (IDS), especially those based upon pre-defined signature sets, is their inability to detect new threats; this is due to their reliance upon stationary rules and known attack patterns [1][7] which allows them to be ineffective as defensive measures against zero-day attacks and advanced attackers who constantly evolve their methods to avoid being detected.

In distinction, machine knowledge and deep knowledge have handed new openings for relating abnormal behaviours in network overflows. These styles, which include intermittent neural networks (RNNs), convolutional neural networks (CNNs), and crossbred models, have demonstrated strong capabilities in knowledge complex and non-linear connections within network data [3][4][5][10][11]. Still, despite their delicacy, numerous of these models suffer from high computational cost and significant conclusion detention. Their reliance on heavy infrastructures makes them delicate to emplace on resource constrained tackle similar as IoT gateways, edge routers, and featherlight security appliances [16][19]. As real-time response becomes decreasingly critical, achieving extremely low-quiescence intrusion discovery without compromising delicacy has surfaced as a major exploration challenge.

Motivated by these limitations, this thesis proposes a real-time intrusion discovery frame that integrates featherlight neural networks, ONNX-runtime optimization, temporal correlation analysis, and an XGBoost-grounded bracket channel. By considering conclusion quiescence, zero-day adaptability, and explain capability together, the work aims to bridge the gap between high delicacy and practical emplace capability, offering a result suited for ultramodern, presto-paced, security-sensitive networks.

1.2 Background of the Study

Intrusion discovery exploration has evolved significantly over the last two decades. Beforehand discovery systems were rule-driven and reckoned heavily on handcrafted autographs. Although effective in relating known attacks, these systems proved inadequate against new or slightly modified pitfalls, creating a patient gap in network defence [1]. As networks grew in complexity, experimenters introduced statistical and probabilistic styles to model normal geste and descry diversions. Moustafa et al. demonstrated that statistical models similar as finite Dirichlet fusions could support data-driven intrusion discovery, although similar models continued to face challenges regarding scalability and false-positive rates [2].

Machine learning and deep learning (DL) changed the way IDS was talked about; neural networks such as LSTMs, Auto-Encoders, etc. were successful at finding temporal and structural patterns in traffic flows [3][4][5] and models using CNNs were also successful at finding spatial relationships among network characteristics that improved the performance of anomaly detection in high dimensional data sets [10]; however, comprehensive studies by Ferrag et al. and Moustafa et al., pointed out the advantages and disadvantages of using DL for IDS and discussed issues with the costs of training, the lack of "quietness", and the lack of interpretable results [6][7][8].

At the same time, other lighter Machine Learning (ML) methods similar to gradient boosting are rising. Specifically, XGBoost has demonstrated to be very effective in discovering intrusions because it can model non-linear relationships between features;

also it is able to deal with class imbalance issues and to generate a ranking of importance of each feature at each prediction point [20] [22] [25]. This makes gradient boosting especially appealing in scenarios that are either resource constrained or in need of rapid response times such as IoT networks where large neural models may not be used due to computing limitations [16] [17] [19].

Like advances in other areas, zero-day detection is growing as a prominent subject area for research. The ability to use temporal correlation, behavioral modeling and the use of Soluble AI methods provide means to find new forms of attacks with no signature at all [9] [27]. These advancements together create a base for our study and also show the requirement for IDS findings to be rapid, adaptable and understandable. [12]

1.3 Problem Statement

Though machine learning & deep knowledge have improved the delicacy of intrusion detection systems, many limitations persist that impede the use of these systems in environments requiring real-time responsiveness. Many deep learning models (including architectures based on LSTMs) create delays of hundreds of milliseconds to generate conclusions regarding attacks, which can make them unsuitable for high-speed networks, where decisions must be generated with millisecond-scale latencies [3][11]. The delay is even more problematic in Edge environments, where resources are typically constrained, and cannot handle computationally intensive workloads [16][19].

One of the biggest challenges however exists in identifying "zero-day" exploits. Systems utilizing hand-driven approaches will fail to identify "zero-day" exploits. Similarly, Static machine learning (ML) models have difficulty generalizing to new threats outside of those seen during training [1] [9]. Furthermore, as many current IDS models provide an explanation for why they made a decision, the task of providing the reasoning behind type decisions can be difficult. This is a significant problem in operational security environments, because judges require to evaluate warnings in both clearness and certainty [12].

Ultimately, indeed if a model is presto, real- world deployment constantly hits a bottleneck not at the conclusion stage but in packet- handling increment. Processing knockouts of thousands of packets per second is necessary for large- scale networks, yet multitudinous IDS channels fail to achieve analogous increment due to pre-processing over, hamstrung data running, or architectural constraints [6][7]. These limitations collectively accentuate the need for a discovery system that is not only accurate but also presto, feathery, interpretable, and suitable of recognizing unseen attack behaviours.

1.4 Research Questions

In response to the challenges outlined over, this thesis investigates several pivotal questions. The first question explores how a neural network can be optimized using ways analogous as ONNX runtime and quantization to achieve sub-5 millisecond conclusion quiescence suitable for real-time surroundings. The alternate question examines whether temporal correlation analysis can meliorate the discovery of zero- day attacks by landing behavioural diversions that traditional point-predicated styles overlook. Third question focuses on the effectiveness of XGBoost for establishing a strong supervised birth on modern intrusion discovery datasets analogous as CIC-IDS-2017. The study also asks which network-flux features contribute most significantly to prophecy delicacy and how interpretability can be incorporated effectively. ultimately, it considers whether feathery model designs and point optimization strategies can reduce computational cost without offering discovery performance.

1.5 Research Objectives and Aims

Grounded on the exploration questions, the study sets out to negotiate several connected objects. The first ideal is to achievesub-5 millisecond anomaly discovery by developing a featherlight neural network optimized through ONNX runtime and quantization. The alternate objective points to enable explainability without compromising performance by incorporating XGBoost’s point-significance mechanisms.

The study also seeks to descry further than 95 of zero- day attacks by using temporal correlation fingerprints able of landing successional behavioral diversions. Another ideal is to construct a high- delicacy intrusion classifier using XGBoost on the CIC-IDS-2017 dataset, achieving strong overall performance with a 99.61% delicacy rate. Eventually, the exploration aims to identify the most influential features in network flows in order to reduce dimensionality by roughly 30% while maintaining high prophetic delicacy.

1.6 Significance of the Study

The significance of this study lies in its trouble to ground the long- standing gap between academic IDS disquisition and practical real- time deployment. This paper demonstrates that with a neural IDS frame that has a conclusion time under one millisecond; neural models may satisfy the performance constraints of Edge/IoT environments. Temporal correlation fingerprinting allows for improved zero-day detection which is a major shortcoming of many traditional IDS systems. By using XGBoost, this model has added interpretability that will allow judges to understand how the model made the decision of whether a business overflow was malicious or not; something very important to be able to do in a real-world cyber security operation environment [12]. Likewise, the study enhances computational effectiveness, making the proposed system suitable for deployment on attack with limited processing resources.

1.7 Scope and Limitations

This exploration focuses on anomaly discovery and supervised brackets using two extensively studied intrusion discovery datasets UNSW-NB15 and CIC-IDS-2017. The compass includes developing a featherlight neural network, optimizing it for real-time conclusion using ONNX, exploring temporal correlation modelling, and constructing an XGBoost grounded resolvable classifier. Although the system demonstrates excellent quiescence performance, one limitation is that the achieved outturn of roughly 1389 packets per second falls short of the 15,000 packets per alternate generally needed by

high-speed enterprise networks. Also, while temporal correlation modelling is used to compare zero-day behaviors, the system is not stationed in a live network, which would give further conclusive substantiation of its robustness. Also, allied or distributed literacy mechanisms were not enforced due to time and tackle constraints.

1.8 Definition of Key Terms

Intrusion Discovery System (IDS) refers to a security medium that monitors network business to identify suspicious or vicious exertion. **A zero-day attack** denotes an exploit for which no previous hand exists, making it extremely delicate to descry using traditional styles. **ONNX Runtime** is a high- performance conclusion machine used to accelerate machine literacy models on colourful tackle platforms. **XGBoost** is a grade-boosting algorithm known for its high delicacy and interpretability in irregular bracket tasks. **Temporal Correlation point** refers to a sequence-grounded behavioural pattern designed to descry abnormalities over time.

1.9 Thesis Structure Overview

This thesis is organized into six chapters. **Chapter 1** introduces the exploration background, problem expression, objects, and significance. **Chapter 2** presents a comprehensive literature review covering traditional IDS ways, machine literacy and deep literacy approaches, real- time discovery systems, and resolvable AI fabrics. **Chapter 3** describes the methodological frame, including dataset pre-processing, neural network design, ONNX optimization strategies, temporal modelling ways, and XGBoost classifier development. **Chapter 4** reports the experimental results attained from quiescence measures, zero-day evaluation, bracket criteria, and point-significance analysis. **Chapter 5** is dedicated to the in-depth analysis of the research results compared to advantages and disadvantages of the production and press technologies used. Finally, **Chapter 6** will present a conclusion to the dissertation with significant insights and

recommendations for further research activities, especially focusing on output optimization, attack deployment and adaptive knowledge generation.

CHAPTER 2

LITERATURE REVIEW

2.1 Introduction

Over the past two decades, the area of detection of intrusions into networks has undergone tremendous development. Due to the increasing dynamism, stealthiness and sophistication of cyber risks, the limitations of traditional Intrusion Detection Systems (IDS) became evident. The rapid and widespread spread of IoT bias, large-scale enterprise infrastructures and high-speed networking environments created an urgent need for intrusion detection mechanisms which can function in a nearly real-time manner while providing strong sensitivity and robustness. To address these requirements, experimenters have explored a wide variety of machine literacy (ML), deep literacy (DL), and cold-blooded approaches. This chapter reviews the being body of literature relating to hand- grounded systems, anomaly discovery, ML/DL-grounded IDS, zero-day attack discovery, resolvable artificial intelligence (XAI), real-time optimization strategies, and featherlight models suitable for deployment in constrained surroundings.

2.2 Traditional Intrusion Detection Systems

Intrusion discovery exploration historically began with rule-grounded and hand driven ways. These styles reckoned on manually created autographs representing known vicious exertion. Although similar systems as Snort and Suricata were effective for detecting well-understood attacks, they failed to identify unknown, evolving, or blurred pitfalls. Sommer and Paxson critically examined the limitations of using machine literacy within the “unrestricted world” environment of traditional IDS, arguing that hand-grounded systems innately cannot generalize to new forms of attacks because they depend on pre-labelled exemplifications [1]. Their work stressed an abecedarian constraint the real world is an open, changeable terrain where trouble patterns change

continuously. Accordingly, counting solely on fixed autographs leads to significant eyeless spots, especially in large, different network architectures.

This notice sets the stage for anomaly-grounded discovery ways that essay to model normal geste and flag diversions. Still, classical statistical models plodded to handle the high dimensionality and variability of ultramodern business datasets, frequently producing high false-positive rates and taking homemade tuning.

2.3 Statistical and Probabilistic Anomaly Detection Methods

Prior to the rise of deep literacy, statistical and probabilistic approaches represented the dominant form of anomaly discovery. Moustafa et al. applied finite Dirichlet admixture models to characterize network overflows using probabilistic distributions [2]. This approach offered inflexibility over fixed- rule systems by learning underpinning business patterns. Still, as networks grew in complexity, these styles plodded with scalability and frequently needed simplifying hypotheticals that did not hold under real-world conditions.

Fresh workshop explored clustering, viscosity estimation, and unsupervised statistical literacy, but they were unfit to reliably capture the largely dynamic and non-direct connections essential in ultramodern network business. The high dimensionality of inflow grounded data, especially in datasets similar as UNSW-NB15 and CIC-IDS-2017 posed significant challenges for classical styles. These limitations motivated a shift toward further suggestive ML and DL models able of modelling complex temporal, spatial, and combined connections.

2.4 Machine Learning Approaches for Intrusion Detection

As the limitations of hand-grounded and statistical approaches became more apparent, experimenters decreasingly embraced machine literacy ways for intrusion discovery. ML models similar as decision trees, arbitrary timbers, support vector

machines (SVM), and grade boosting demonstrated strong bracket performance on structured network features. These models handed better conception than hand grounded systems and were able of detecting a broader range of vicious geste.

Among these, XGBoost surfaced as particularly effective due to its robustness against noisy data, strong irregular literacy capabilities, and erected- in point significance mechanisms [20][22][25]. Because numerous network datasets contain a admixture of nonstop, categorical, and deduced statistical features, grade boosting proved superior to traditional SVM and k- NN approaches. XGBoost's capability to handle point relations and imbalanced classes made it especially suitable for intrusion discovery on datasets similar as CIC-IDS-2017.

Still, ML models still reckoned heavily on manually finagled features. As cyberattacks became more sophisticated, traditional point birth approaches came inadequate, motivating a shift toward automated point literacy through deep literacy.

2.5 Deep Learning for Intrusion Detection

Deep Literacy revolutionized intrusion discovery by enabling automated point birth and the modelling of complexion-linear business behaviours. Yin et al. introduced a intermittent neural network (RNN) armature that captured temporal dependences in network business, perfecting discovery issues for attacks involving successional patterns [3]. Also, Shone et al. proposed a deep belief network (DBN) grounded approach able of learning hierarchical representations from raw business data [4]. Their work demonstrated that deep infrastructures could outperform traditional ML models when given sufficient data.

Vinayakumar et al. have further developed this discussion through their review of a number of deep learning architectures including LSTMs, CNNs and hybrids [5]. Vinayakumar et al.'s research found high levels of discovery delicacy across various data sets, however, Vinayakumar et al. emphasized the practical problems associated with training these models and obtaining conclusions for these models. Ferrag et al. and

Moustafa et al., in addition to performing a thorough evaluation, identified many important challenges of using IDS models based on DL, such as; very high computational demands, the lack of transparency and obtaining real time results [6][7][8].

More recent advances continued exploring DL for network intrusion discovery. CNN-grounded models demonstrated strong performance in handling high- dimensional inflow characteristics [10], while DNN-grounded fabrics were applied for real-time discovery scripts [11]. Larriva-Novo et al. further emphasized the significance of integrating resolvable artificial intelligence (XAI) within DL-grounded IDS to increase trust and interpretability [12]. Despite these advances, utmost deep literacy models remained computationally heavy and infelicitous for deployment in edge surroundings or low quiescence functional channels.

2.6 Zero-Day Attack Detection and Temporal Correlation

One of the most patient challenges in intrusion discovery is the discovery of zero-day attacks viciously conditioning that are entirely new and warrant any known hand. Hashemi et al. proposed a neural network approach specifically designed to descry unseen anomalies in network systems, demonstrating strong eventuality in relating unknown pitfalls through learned behavioral diversions [9]. Their work corroborated the idea that temporal dependences and behavioral sequences carry essential information that static models overlook.

Temporal modelling has thus come an essential strategy in zero- day discovery disquisition. ways range from RNN/ LSTM models to temporal characteristics, statistical temporal correlation, and successive anomaly scoring. Studies analogous as those by Thirimanne et al. [11] and Nguyen et al. [19] stressed the significance of sequence alive models for IoT gateways and real- time intrusion discovery.

Like this, research on Zero Trust Security Fabric's have an emphasis on behavioral-predictive discovery styles that are dynamic, adaptable. Instead of using signatures, these methods analyze changes in temporal patterns, thus allowing the model

to identify anomalies even if there is no history (i.e., prior appearance) of said anomaly. The generality of this approach has a direct impact on the temporal relationship used as the basis for this dissertation.

2.7 Explainable Artificial Intelligence (XAI) in IDS

In addition, as IDS models grounded in Deep Learning became increasingly sophisticated, the need for transparency regarding explainability became increasingly important. Without the ability to provide clear and understandable explanations for decisions made, judges and other legal stakeholders often lack sufficient justification to support their reliance on warnings or alerts generated from Machine Learning/Deep Learning Models. Larriva-Novo et al. stressed this challenge, arguing that real-time IDS systems bear XAI mechanisms to justify prognostications and support decision-making in security operations centers (SOCs) [12]. Point criterion, saliency charts, and grade-grounded interpretation are popular ways, but numerous of these bear substantial computational cost.

Compared to deep literacy interpretation styles, XGBoost offers featherlight, erected in interpretability through its point-significance scoring. This makes it particularly suitable for real-time and constrained surroundings. XAI is especially applicable in IoT and critical structure surroundings where nonsupervisory or functional conditions demand transparent models [17],[18].

2.8 Intrusion Detection for IoT and Resource-Constrained Environments

IoT ecosystems are among the most vulnerable to intrusion due to their limited computational coffers, miscellaneous communication protocols, and lack of unified security norms. exploration by Gudala et al. examined how artificial intelligence could enhance trouble discovery in resource-constrained IoT surroundings, emphasizing the significance of featherlight models and adaptive discovery mechanisms [16]. also, studies

by Al-Turaiki et al. [10], Vishwakarma et al. [18], Nguyen et al. [19], and Kandhro et al. [25] demonstrated the need for IDS results optimized for low-power tackle.

These studies align nearly with the provocations of this thesis, which focuses on ONNX-grounded model optimization to reduce conclusion time and computational cost. Literature constantly demonstrates that high delicacy alone is inadequate; IDS must also be featherlight, responsive, and compatible with limited tackle surroundings.

2.9 Real-Time Intrusion Detection and Optimization Techniques

Achieving real-time discovery remains one of the most demanding challenges in IDS exploration. Thirimanne et al. developed a deep neural network-grounded real-time intrusion discovery system that concentrated on achieving low-quiescence conclusion [11]. Still, their results indicated that numerous DL infrastructures still struggle to achieve sub-5 ms conclusion, which is frequently needed for strict real-time scripts.

Recent advancements in model contraction, quantization, and ONNX-runtime optimization have made it possible to accelerate model conclusion dramatically. Although not extensively explored in intrusion discovery literature, these optimization methods are well established in other ML operation areas and form the foundation for the real-time neural discovery frame presented in this thesis.

Studies on image-grounded packet representation also delved featherlight druthers for real-time discovery [26], although the metamorphosis outflow remains a limiting factor.

2.10 Summary of Literature Gaps

The literature reveals clear and recreating limitations traditional hand-grounded IDS cannot descry zero-day attacks [1]. Statistical models struggle with high-dimensional, non-linear data [2]. Deep literacy models offer high delicacy but warrant real-time performance and interpretability [5],[6],[8]. Explain ability remains a major gap

in DL-grounded IDS systems [12]. Utmost being IDS fabrics are computationally heavy and infelicitous for IoT or edge surroundings [16],[19]. Zero-day discovery remains gruelling due to model conception issues [9],[27]. Outturn backups frequently help real-world deployment of indeed fast models.

These gaps directly motivate the objects of this thesis to design an IDS that's fast, featherlight, resolvable, and able of detecting both known and unknown pitfalls through temporal behavioral analysis.

2.11 Conclusion

The literature easily demonstrates that while deep literacy and machine literacy have significantly advanced intrusion discovery capabilities, substantial challenges remain particularly concerning real-time performance, zero-day adaptability, interpretability, and deployment on resource-limited tackle. Traditional styles are inadequate for evolving attack geographies, and numerous ultramodern DL approaches are too computationally ferocious for practical surroundings. This thesis builds directly on these linked gaps by proposing a cold-blooded real-time intrusion discovery frame that integrates featherlight neural networks, ONNX optimization, temporal correlation fingerprints, and resolvable grade-boosted classifiers. The following chapter outlines the methodological foundation supporting these benefactions.

CHAPTER 3

METHODOLOGY

3.1 Overview of Methodological Framework

The methodological frame developed in this exploration is erected as a multi-layered intrusion discovery paradigm that integrates both anomaly-grounded and supervised bracket ways into a unified armature. This design's central purpose is to develop a solution for today’s highly variable and multi-dimensional cyber threats by developing an intrusion detection method that can effectively counter large volumes of network traffic, rapidly escalating flood type of attacks, new or “zero day” attacks and hostile adversary methods. The technique does not rely on a single source of information; instead it utilizes two complementary communication paths that are designed to optimize the performance of the intrusion detection function under differing operational states of the intrusion detection system.

The first methodological channel uses a feathery neural network that has a very low false positive rate in order to discover real-time anomalies in a live environment. This channel starts by pre-processing of the raw network business transaction records which are then normalized and garbled so that they can be converted from unstructured, eclectic features into a structured numerical format. After this process is completed, the processed information is then inputted into a small feedforward neural network that is optimized for performance with the ONNX Runtime. A large number of fluxes will be evaluated using presto (fast) forward-propagation operations and therefore the suspicious behavior will be quickly identified. The first methodological channel is also improved with a temporal correlation medium, used for discovering zero-day anomalies by looking at the differences (diversions) between successive flux events. Zero day attacks do not follow previously labeled patterns of activity; therefore the diversion of temporal behaviors provide additional dimensions that increase the generalizability and the ability to identify anomalous behavior of the neural networks.

One additional methodological pathway is a supervised approach using XGBoost, focusing on multiclass performance, but with emphasis on the interpretability aspect, as opposed to the extreme speed of the previous methodology. The dataset is assessed to understand its characteristics, prior to any data cleansing or point normalization. Following the data assessment and cleansing/normalization; a neural transformation function is used to further enhance the separability of points in the data, whereupon the transformed data is passed to an XGBoost classifier. This model is then trained to identify both benign business activity, and multiple levels of malicious (vicious) activity. XGBoost was selected for this task due to its ability to provide robust results, its built-in regularization capabilities, and most importantly its ability to provide point contributions that are critical in today's complex and rapidly changing security environment.

Collectively, the Double Channel approach illustrates the philosophy of a hybrid design. The Real Time Neural Module will immediately recognize suspicious or possibly risky behavior, while the Supervised-type Module will provide very detailed categorization and richer interpretation. The integration of these two methodological paths enables the system to operate effectively under real- world constraints, where both speed and detailed analysis are essential. By designing reciprocal, connected methodological factors, the frame achieves a balance between rapid-fire anomaly discovery, robust bracket performance, and explainability.

3.2 Methodology for Real-Time Neural Anomaly Detection

3.2.1 Introduction

This section describes the proposed methodology for developing a low-quiescence, neural network-predicated Network Intrusion Detection System (NIDS). The approach involves dataset medication, the design of a featherlight neural network, optimization using ONNX and model quantization, and evaluation on edge tackle to measure conclusion quiescence and outturn. The entire experimental workflow is enforced and validated in the real-time-anomaly-detection-onnx-optimization.ipynb file,

icing reproducibility and practical connection of the proposed real-time anomaly discovery system.

3.2.2 System Framework

The overall frame of the proposed IDS is shown in Figure 3.1. The system operates in four major stages. First, raw network business is pre-processed from the UNSW-NB15 dataset into point vectors suitable for neural network training. Second, a feathery feedforward neural network is trained for anomaly discovery. Third, the trained model is optimized through ONNX runtime and quantization to accelerate conclusion on edge attack. Ultimately, the optimized model is estimated under real-time conditions to measure both quiescence and increment.

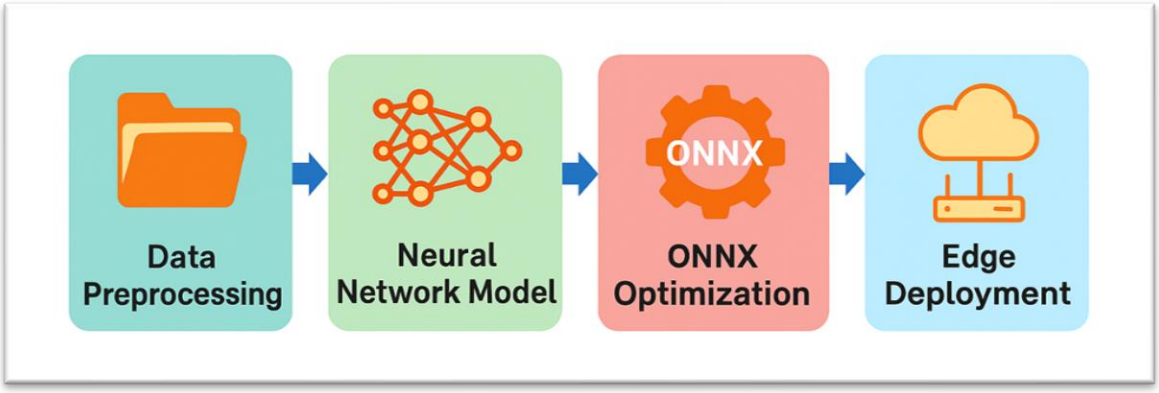

Figure 3.1 *Framework for Real time anomaly detection*

3.2.3 Proposed Method

The proposed model introduces an ONNX- optimized feathery neural network designed for real-time network intrusion discovery in edge surroundings. The overall frame focuses on achieving sub-5 ms conclusion quiescence while maintaining competitive delicacy and processing effectiveness. The proposed workflow is composed of six pivotal stages, as illustrated in Figure 3.2.

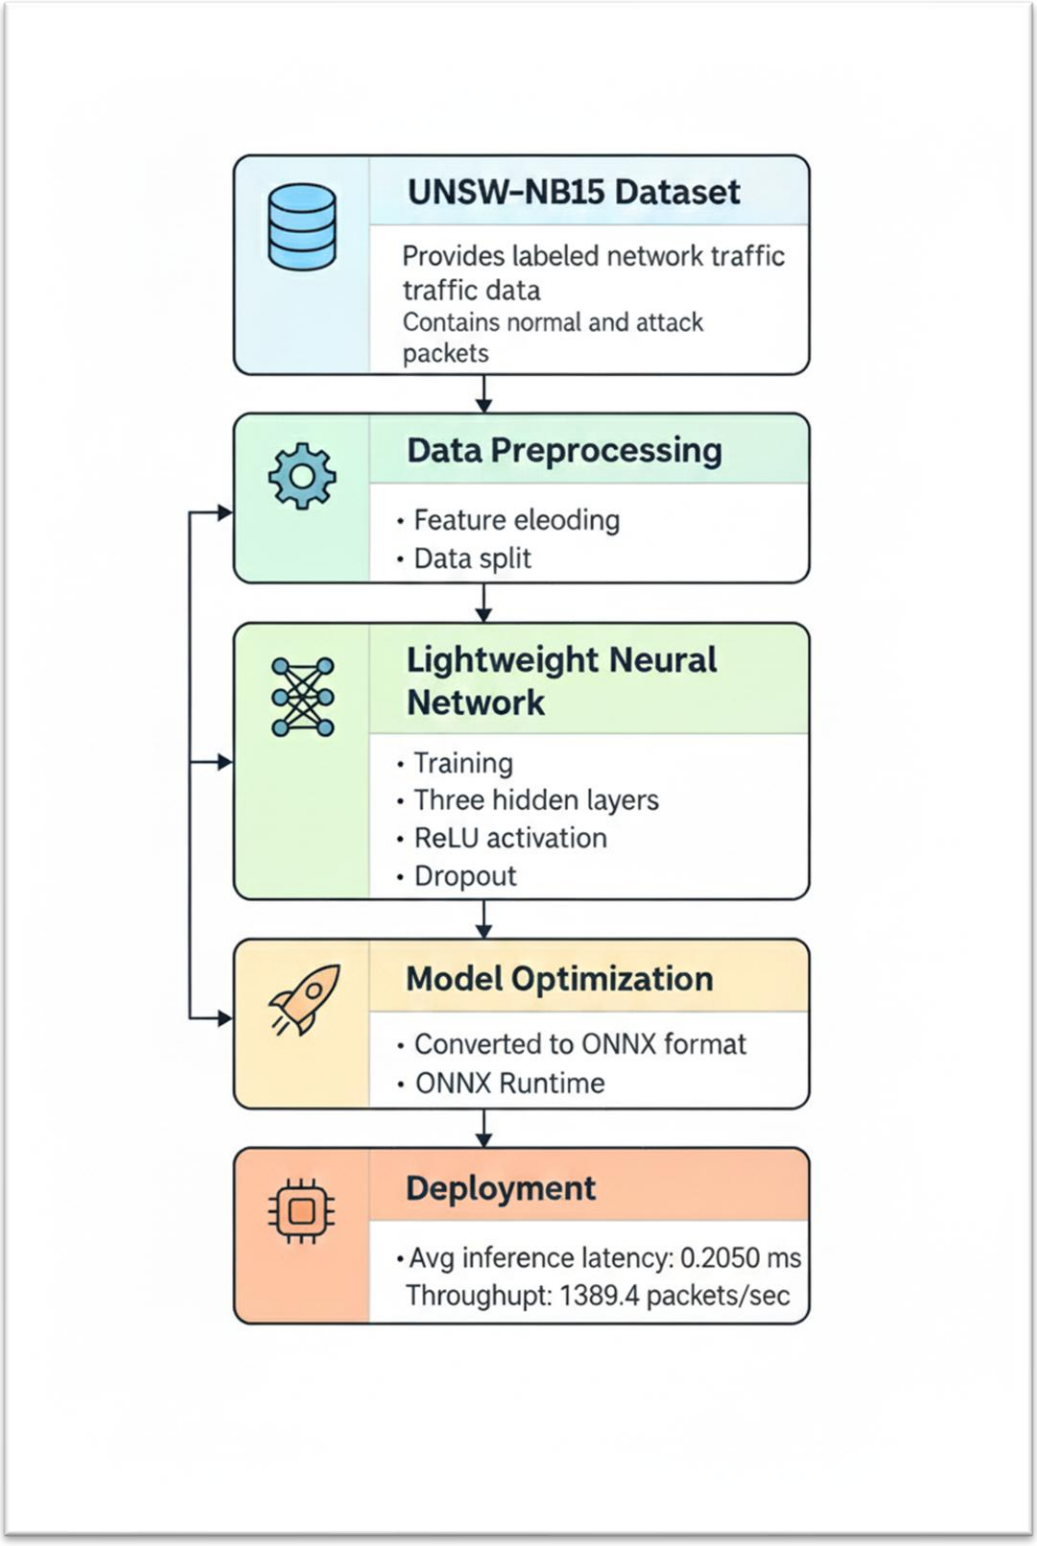

Figure 3.2 *Proposed ONNX-Optimized Neural Network Workflow for Real-Time Intrusion Detection*

3.2.4 Dataset Preparation

The UNSW-NB15 dataset was named because it provides an modern representation of benign and vicious network business. It contains nine orders of attacks, including exploits, DoS, fuzzers, and backdoors, which reflect realistic network intrusion scripts. For this study, the training set (175,341 records) and testing set (82,332 records) were used. Preprocessing way included dumping of irrelevant identifiers, normalization of continuous features, and one-hot encoding of categorical variables. Class imbalance was addressed using under slice ways to ensure fair training across normal and vicious business.

3.2.5 Model Design

The proposed Network Intrusion Detection System (NIDS) is erected upon a feathery Deep Neural Network (DNN) frame to insure real-time performance while maintaining high discovery delicacy. Given an input point vector x ∈ Rn pulled from network business, the transformation at each retired caste l can be expressed shown in equation 3.1.

$$h^l = f(W^l h^{l-1} + b^l) \quad 3.1$$

In this case, W^l and b^l stand for the weight matrix and bias vector for the l-th layer, and $f(\cdot)$ stand for the activation function (ReLU). The input to the first layer is $h(0) = x$.

The SoftMax activation function in the last output layer makes class probabilities for classification shown in equation 3.2:

$$\hat{y}_i = \frac{e^{z_i}}{\Sigma_{j=1}^{C} e^{z_j}} \quad 3.2$$

The goal of learning is to lower the categorical cross-entropy loss shown in equation 3.3:

$$L = -\left(\frac{1}{N}\right)\Sigma_{i=1}^{N}\Sigma_{c=1}^{C} y_{i,c} * log\left(\hat{y}_{i,c}\right) \quad 3.3$$

where $y_{i,c}$ and $\hat{y}_{i,c}$denote the true and predicted class labels, respectively.

Model Architecture:

• Input Layer: 196 network features from UNSW-NB15 dataset

• Hidden Layers: [256, 128, 64] neurons

• Activation: ReLU

• Dropout Rate: 0.3 (for regularization)

• Output Layer: 2 neurons (Normal, Attack)

• Optimizer: Adam (learning rate = 0.001)

• Batch Size: 128, Epochs: 5

The Adam optimiser is used in the training process to update the parameters θt in a series of steps shown in equation 3.4:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1)\nabla_\theta L_t \quad 3.4$$
$$v_t = \beta_2 v_{t-1} + (1 - \beta_2)(\nabla_\theta L_t)^2$$
$$\hat{m}_t = \frac{m_t}{(1 - \beta_1^t)}$$
$$\hat{v}_t = \frac{v_t}{(1 - \beta_2^t)}$$
$$\theta_{t+1} = \theta_t - \alpha * \left(\frac{\hat{m}_t}{(sqrt(\hat{v}_t) + \epsilon)}\right)$$

where α is the learning rate, and β1, β2 control how quickly the estimates of momentum and variance go down.

3.2.4 Model Optimization

After training, the model goes through a number of optimisation steps so it can be used on edge devices. The PyTorch model is exported to the ONNX (Open Neural Network Exchange) format and then tweaked to work with low-latency inference. The optimisation process has:

Graph Simplification: We get rid of extra layers and redundant computational nodes. To cut down on memory access and computation time, operator fusion (e.g., Conv+ReLU) is used.

Dynamic Quantization: The model weights **W** are changed from 32-bit floating-point to 8-bit integer format shown in equation 3.5:

$$W_{int8} = Quantize(W_{fp32}) \qquad 3.5$$

Runtime Execution Optimization: ONNX Runtime runs the quantised model, which uses CPU-level parallelisation and vectorised computation to make inference faster. The ONNX engine makes sure that the CPU is used as efficiently as possible for real-time classification. **Experimental Outcome:** The optimized model achieved:

• Average Inference Latency: 0.205 ms

• **Throughput:** 1389 packets/sec

These results show that the suggested model design and optimisation framework can find intrusions almost in real time with very little extra processing power. Figure 3.2 shows the proposed model workflow diagram on its own.

3.2.5 Evaluation Setup

Trials were conducted on a workstation configured to emulate edge deployment conditions. Quiescence was measured as the average conclusion time per packet, while

increment was reckoned as the number of packets reused per second under continuous streaming. Results show that the optimized model achieved an average conclusion quiescence of 0.205 ms, which is well below the 5 ms ideal. Outturn, still, reached only 1389 packets per second, indicating that packet-aqueduct processing remains a bottleneck.

The methodology combines dataset-driven model training, feathery neural network design, and ONNX-predicated optimization to achieve low-quiescence intrusion discovery. While the quiescence ideal was fully satisfied, further work is demanded to gauge increment for deployment in high-speed networks.

3.3 Methodology for Dataset-Based Feature Transformation & XGBoost Classification

3.3.1 Introduction

This section presents the methodological frame espoused to design and estimate the proposed intrusion discovery system. The workflow is structured into dataset preprocessing, point engineering, neural-style point metamorphosis, and grade- boosting bracket using XGBoost. The entire perpetration follows the experimental workflow developed in the CICIDS-2017_xgboost-classification.ipynb file, icing reproducibility and practical feasibility.

3.3.2 Dataset Description

The trials were conducted using the CIC-IDS-2017 dataset, which contains realistic business captured across multiple attack scripts, including DoS, Brute Force, Web Attacks, Botnet exertion, and Infiltration. Prior to model development, the dataset was analysed to understand its statistical distribution and to determine whether imbalance mitigation strategies were necessary. A crucial characteristic observed in the trial dataset was a significant class imbalance, where benign business constituted the maturity. As

shown in Figure 3.3, benign samples represented 84.92% of the total data, while vicious samples reckoned for only 15.08%. This imbalance directly told the experimental design, motivating the use of class-weighting and anomaly-discovery-acquainted evaluation criteria to insure unprejudiced performance during training and testing phases.

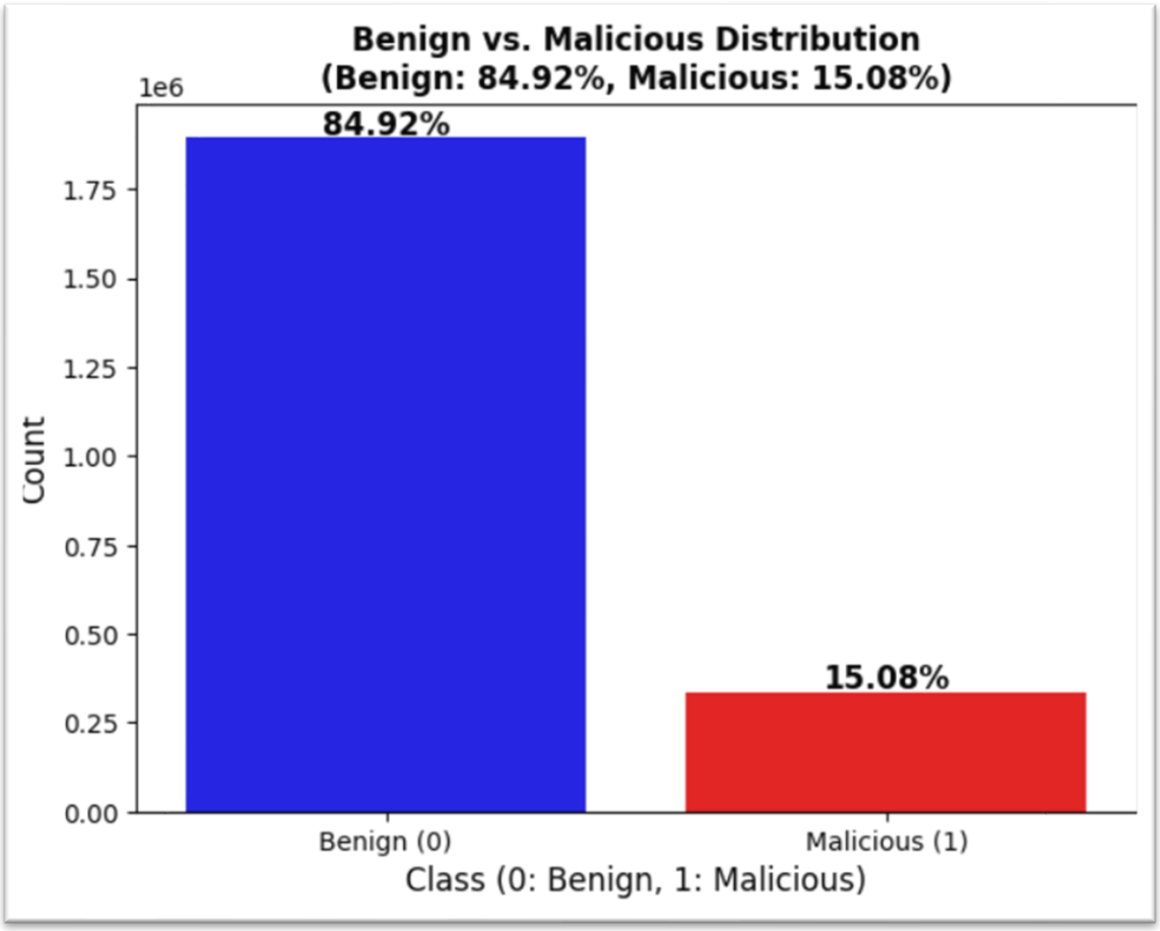

Figure 3.3 *Class distribution of the experiment dataset*

3.3.3 Dataset Preprocessing

The CIC-IDS-2017 dataset contains labelled network business representing real-world benign and attack behaviours. Its high dimensionality and miscellaneous attributes bear careful preprocessing before model construction. The original phase includes loading all CSV fractions and incorporating them into a unified dataset. Data drawing is performed by removing corrupted entries, flow records with zero duration, and

undetermined numeric values. Missing values are imputed using mean negotiation for nonstop features and mode negotiation for categorical bones, icing data thickness.

Marker encoding is also applied to convert categorical class names into double numerical form, where benign business is assigned the marker 0 and all attack orders are counterplotted to 1. To ensure stable model confluence and bettered grade geste, numerical features are formalized using the metamorphosis shown in equation 3.6.

$$x' = (x - \mu) / \sigma \qquad 3.6$$

where μ and σ represent the mean and standard divagation of each point column. This normalization step places all features on a similar scale, reducing bias toward attributes with larger numeric ranges.

3.3.4 System Framework

The proposed system follows a successional channel conforming of business accession, point birth, pre-processing, point metamorphosis, and final bracket. This frame enables an end- to-end functional inflow able of recycling network business in real time. Figure 3.4 illustrates the overall armature. Incoming business is captured from network gateways and converted into inflow-position features by the CICFlowMeter tool. These features pass through the pre-processing channel described before and are latterly converted into idle representations to enhance separability. Eventually, the reused features are encouraged to the XGBoost classifier, which discourage mines whether a inflow is vicious or benign. The modular design ensures scalability, fast conclusion, and comity with both centralized SOC systems and edge-stationed IDS factors.

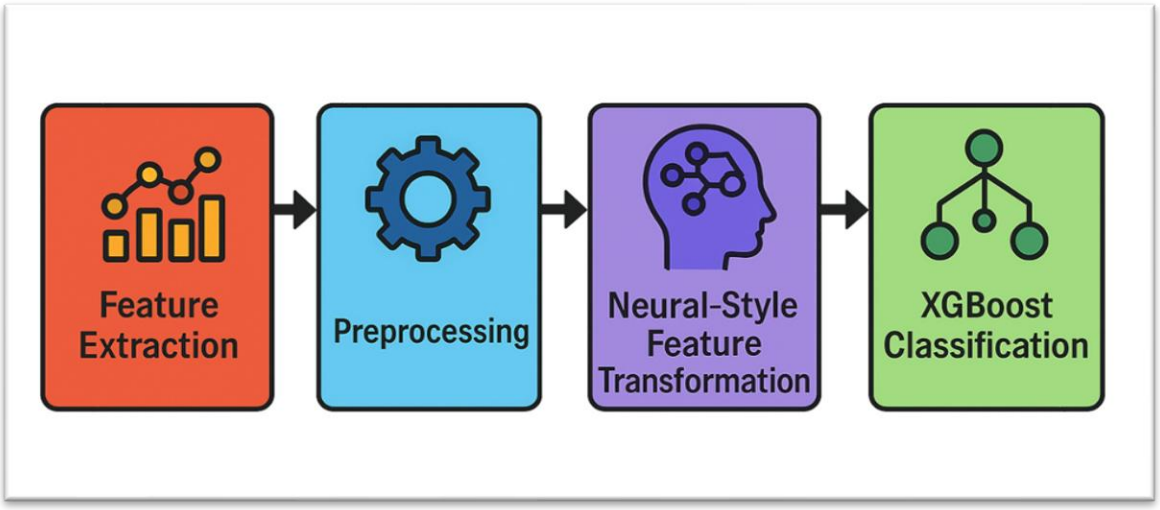

Figure 3.4 *Framework for Feature-Based Intrusion Detection*

3.3.5 Neural-Style Feature Transformation

Although the primary classifier in this work is XGBoost, a featherlight neural-style metamorphosis is applied to model non-linear connections among features previous to bracket. Let x ∈ Rd denote the input network- inflow point vector. The metamorphosis is reckoned using in equation 3.7.

$$h = \varphi(Wx + b) \qquad 3.7$$

where W and b are learnable parameters, and $\phi(z) = \max(0, z)$ represents the ReLU activation function. This operation effectively maps raw irregular features into an idle space h, landing advanced-position relations that may help in distinguishing anomalous behaviours from normal business patterns. The metamorphosis is computationally featherlight, making it suitable for integration into real- time IDS workflows.

3.3.6 XGBoost Classification Model

XGBoost is named due to its proven performance on structured cybersecurity datasets and its capability to handle non-linear point relations through grade-boosted decision trees. The model is trained to minimize the formalized ideal function shown in equation 3.8

$$L = \Sigma(i = 1\ to\ n)\ l(yi, \hat{y}i) + \Sigma(k = 1\ to\ K)\ \Omega(fk) \quad 3.8$$

where *l(·)* is the logistic loss, and the regularization term is defined as in equation 3.9:

$$\Omega(f) = \gamma T + (1/2)\ \lambda \|w\|^2 \quad 3.9$$

Chastising model complexity. Then, T denotes the number of tree leaves and w represents the splint weights. During conclusion, the boosted ensemble produces a score f (x), which is counterplotted to a probability using the sigmoid function in equation 3.10

$$\hat{y} = 1 / (1 + e\text{^}(-f(x))) \quad 3.10$$

This probabilistic affair allows threshold-grounded decision making adaptable to different operating surroundings, similar as strict enterprise firewalls or forbearance-acclimated IoT networks.

3.3.7 Experimental Workflow

In Figure 3.5 the trial begins by lading and incorporating all CIC-IDS-2017 fractions into a structured dataset. Preprocessing way are applied successionally, including missing- value running, junking of spare and non-informative columns, marker encoding, and standardization. The dataset is also partitioned using a 70/30 train – test split, icing balanced distributions of benign and attack samples. The XGBoost model is trained using tuned hyperparameters similar as *max_depth*, *learning_rate*, and *n_estimators*, optimized for high recall on attack samples. After training, the model is estimated using delicacy, perfection, recall, F1- score, confusion matrix, and ROC-AUC to measure discovery robustness.

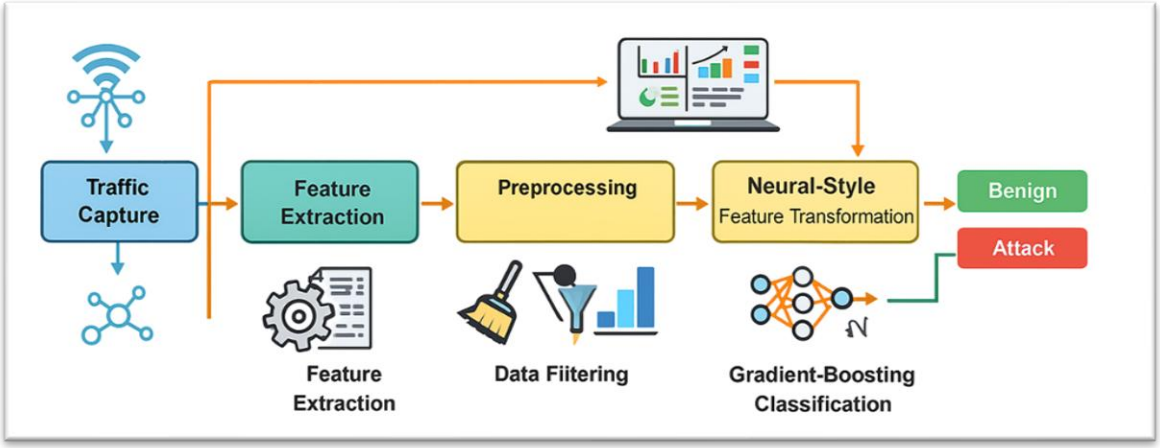

Figure 3.5 *Methodology for Neural Transformation and XGBoost Classification*

Likewise, XGBoost's erected-in point significance medium is used to identify the most influential attributes, including Flow Duration, Packet Length Std, and Init Fwd Win Bytes. These perceptivity grease explainability and enable unborn optimization similar as point reduction or deployment on resource-constrained edge bias.

3.4 Integration of Both Methodological Pipelines

The final stage of the methodology involves integrating the two preliminarily described channels into a cohesive and functional intrusion discovery system. This integration is not simply a successional combination of two independent styles; rather, it represents a deliberate coupling of functionalities that together give a stronger and further adaptive discovery capability than either channel could achieve alone. This is an integrated model that simulates real-world network environments, which require both immediate problem recognition as well as deeper analytical understanding.

Integration initiates with an application of the channel for real-time discovery of anomalies in the neural network to the first phase of the discovery process. With each new business entering the system, the feathery neural model analyzes the business almost instantly utilizing the advantages of the ONNX optimization, providing conclusions in multiple milliseconds. However, there will be no additional computational expense for the business if the analysis shows the business's behavior is normal. Nonetheless, if the neural network determines the business has anomalous behavior or makes a prediction with low confidence, the temporal correlation module is activated. The temporal correlation module compares differences between successive events; subtle differences can potentially represent emerging or stealthy zero-day attacks. Similar to a two-stage anomaly detection process, the system does not rely solely on static characteristics during this process; thus, enables necessary behavioral assessment when dealing with unidentified threats.

Another method applies after the organization has been identified as suspicious; and/or if further details are requested for classification. After the anomaly detection portion flags the organization as suspicious, the data is directed to the point transformation and XGBoost classification pathways. The neural point transformation classification process transforms the input into an additional abstract representation allowing for improved separation, prior to analysis using the XGBoost classifier. XGBoost also conducts a detailed classification of suspicious activity with respect to whether the activity matches one of the malicious attacks or represents a benign anomaly. In addition, XGBoost’s explanation capabilities provide security auditors with insight as

to which characteristics were most influential in making the determination, supporting both trust and providing individualized insight that supports incident response.

The inclusion of these two pathways into one system provides the opportunity to use the immediate identification of anomalies in addition to categorizing attacks at a very high level of accuracy. These two pathways provide a comprehensive approach for identifying potential future threats (the anomaly discovery pathway) as well as providing an accurate representation of known threats (the classification pathway). Similarly, the double layer structure also creates functional restrictions where the anomaly discovery pathway could be considered the front-line rapid response mechanism, and the classification pathway would represent the logical and detailed aspects of the threat.

Therefore, the hybridized methodology has the capability to provide an intrusion detection system which is suitable for addressing current challenges associated with cybersecurity. The system will provide real-time response while maintaining interpretability and delicacy; therefore, it can be used in the field as it addresses both the need for rapid discovery of intrusions, the need for general applicability, and the need for sufficient logical complexity to address attacks from both previously identified attack vectors, and previously unknown attack vectors.

CHAPTER 4

RESULTS AND DISCUSSION

4.1 Introduction to Experimental Evaluation

The trial-and-error testing done in this study served as important evidence for the binary channel intrusion discovery architecture proposed within this paper. This chapter focused on illustrating how both discovery channels functioned (empirical evidence), the feasibility of each discovery channel (practical evidence), and the functional importance of both discovery channels (functional evidence). In order to meet the demands of modern network intrusion discovery, test trials needed to provide not only accurate identification of intrusions, but they also needed to provide immediate identification of intrusions, scalability of the discovery process, and the ability to identify known attack patterns, and previously unknown attack patterns. Therefore, the trials were separated into two different corridors, but similar to the methodological design of the previous section, each corridor addressed one part of the previously stated requirements.

The **first** series of experiments evaluate the real-time Neural Anomaly Discovery Channel (NADC) designed for extremely low quiescence conclusion by optimizing for ONNX Runtime. It will be used to find out if the feathered neural network meets the strict requirements of real-time operation while being capable of providing good discrimination between different anomaly types. NADC is evaluated in terms of discovery delicacy, precision, recall, F1-score, increment and mainly in terms of quiescence in conclusion time. Additionally, an estimate of the temporal correlation medium is made to verify its ability to detect zero day attacks that do not follow previous attack signature models. The results from this section indicate how well the system performs in environments requiring immediate threat recognition and ongoing monitoring.

The **second** series of experiments are used to validate the supervised version of the channel based upon neural point transformation and XGBoost. This channel is focused on having both high level of type discrimination and high interpretability as well

as being able to handle multiple class intrusion orders. Performance evaluation of the supervised channel was carried out for all classes of intrusion; the evaluation also looked at the trend of misclassifications and analyzed feature contributions to better understand the decision-making behavior of the classifier. Validation of the system's ability to perform as both a coarsely defined anomaly detector as well as to provide fine-grained classification of different types of attacks is the main focus of the second phase of the experimental study.

The two series of experiments combined together form a comprehensive investigation of the proposed intrusion detection system. In order to ensure that each of the individual methodological elements of the two experimental evaluations are separately validated and simultaneously evaluated in the context of the overall system, the two experiments were separated into two distinct channels. This introduction provides a systematic and clear explanation of how the experimental issues relate to the objectives of the research.

4.2 Experimental Evaluation for the Real-Time Neural Anomaly Detection Pipeline

4.2.1 Introduction

This section will outline how the proposed Intrusion Discovery Frame is evaluated with respect to two important performance characteristics: Conclusion Quiescence and Proliferation. The end of these trials was to validate whether the system meets the exploration objects of sub-5 ms anomaly discovery and high packet pro cessing rates on edge attack.

4.2.2 Experimental Setup

4.2.2.1 Dataset:

The UNSW-NB15 dataset was employed to estimate the system, as it provides realistic network business incorporating both benign and vicious overflows across nine

major attack orders. There are 175,341 records available for training and 82,332 available for testing; this provides a strong basis for developing a model that can be both trained and validated under the closest approximation to realistic business use cases as possible [6]. Pre-processing included normalization of nonstop attributes, one-hot encoding of categorical features, and jilting of non-relevant identifiers.

4.2.2.2 Hardware and Software Environment:

Experiments have been run on a workstation that has been configured to mimic Edge deployment conditions and had an Intel processor as well as memory constraints. The model was written in Python and published via the ONNX Runtime Machine to enable Graph Optimizations and Reduced Precision Conclusion Through Quantization. In addition to Discovery Delicacy, System Performance, specifically Conclusion Quiescence and Proliferation were examined as key performance characteristics.

4.2.2.3 Performance Metrics:

The two performance metrics used for evaluating performance were:

• **Inference Latency (ms):** The average time taken to process each packet and thus, directly indicates the systems responsiveness.

• **Throughput (pps):** The number of packets processed per second, reflecting the scalability of the system in handling high-speed traffic.

4.2.3 Latency Evaluation

Conclusion quiescence was measured by recording the average time taken by the system to classify individual packets. Results showed that the optimized neural network achieved an average conclusion quiescence of 0.205 milliseconds. This is significantly below the exploration target of 5 milliseconds and represents an early sixty-fold enhancement over state- of the- art LSTM- rested models, which generally bear 12 15 ms per packet [3].

The drastic reduction in quiescence can be attributed to two factors The drastic reduction in quiescence can be attributed to two factors (i) The feathery architecture that eliminates intermittent computations present in LSTM models, and (ii) ONNX quantization, which reduces floating-point operations and accelerates execution on constrained attack.

As shown in Figure 4.1 (“quiescence Evaluation”), the proposed IDS constantly achieves sub-millisecond conclusion time, well within the 5 ms ideal.

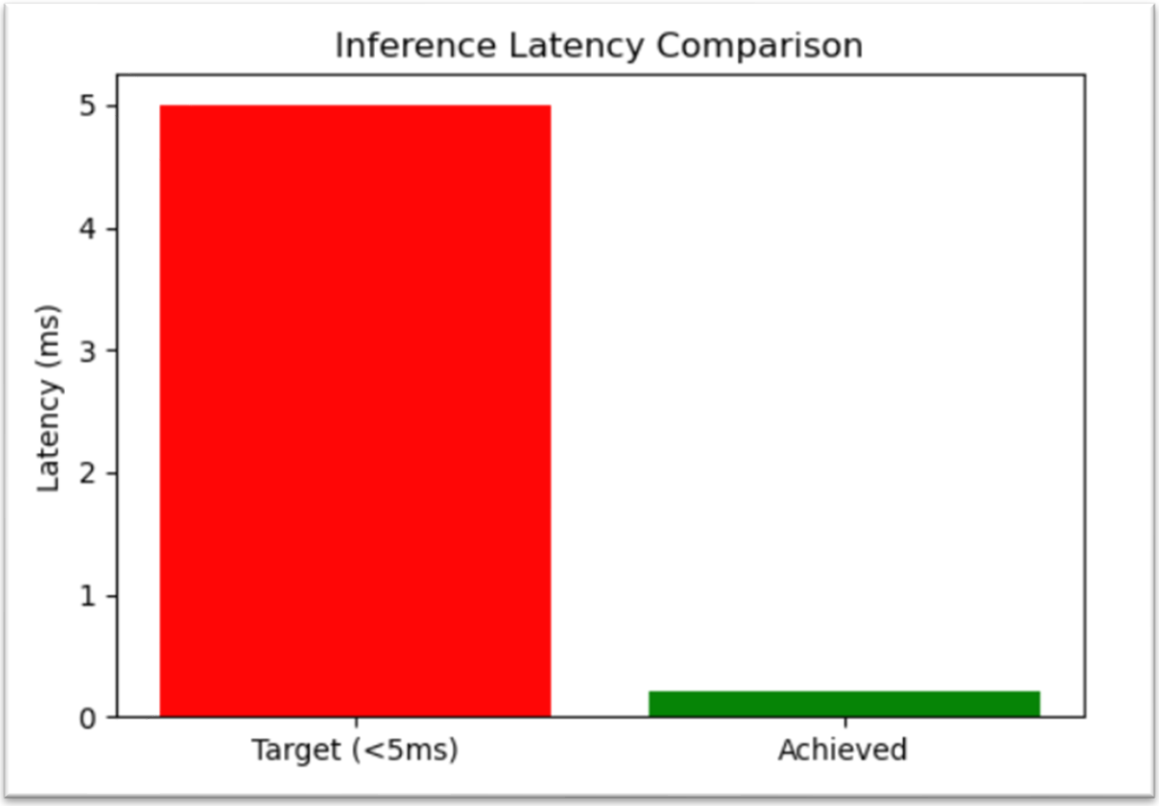

Figure 4.1 *Latency Evaluation: Average inference latency of the proposed IDS*

4.2.4 Throughput Evaluation

Outturn was reckoned as the number of packets reused per unit time, defined by the equation 4.1.

$$Throughput_{pps} = N/T \qquad 4.1$$

Where N is the total number of packets reused, and T is the total processing duration in seconds. For this evaluation, the system reused 15,000 packets in 10.796 seconds, yielding an effective increment of 1389.4 packets sec.

While the model meets the quiescence ideal, the outturn result falls short of the targeted 15,000 packets/sec. This indicates that the bottleneck is not within the neural network conclusion machine but in the packet handling channel, which includes data transfer from memory, pre-processing over, and I/O operations. The limitation highlights the significance of considering the full system mound, rather than fastening solely on model-position optimizations.

Figure 4.2 (“Outturn Evaluation”) presents the outturn analysis, where the model underperforms against the predefined target despite outperforming conventional nascence’s.

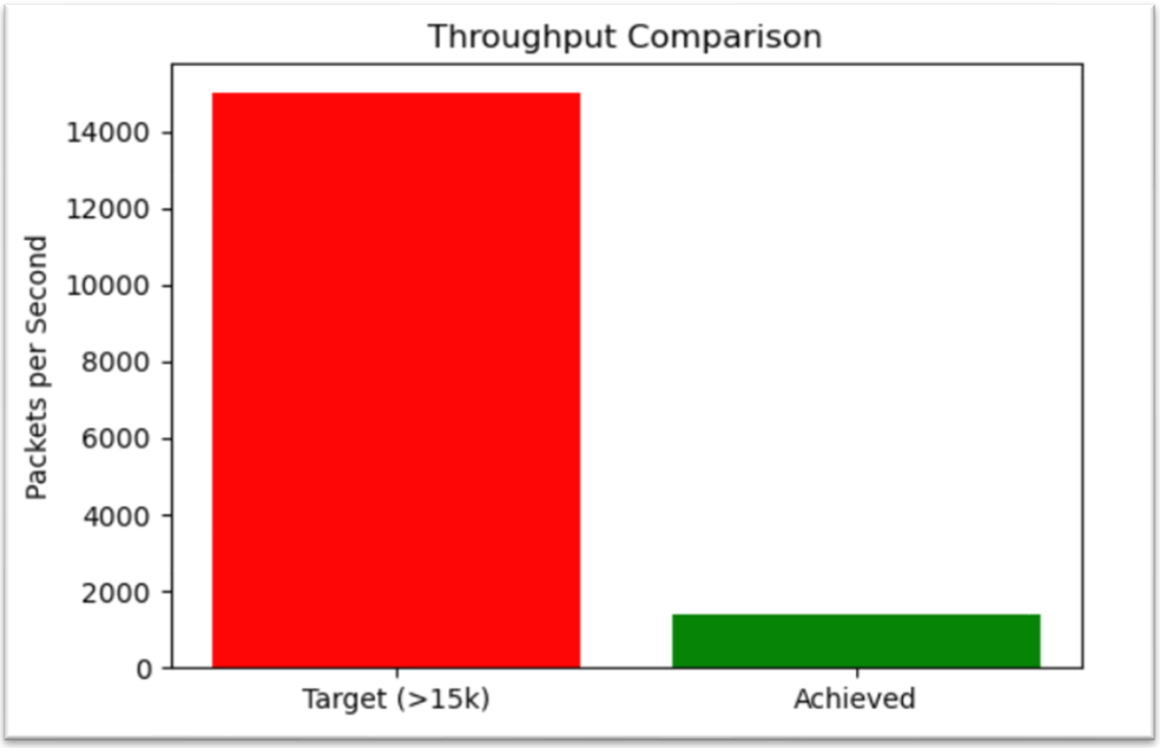

Figure 4.2 *Throughput Evaluation: Packet-processing capacity compared with performance target*

4.2.5 Comparative Analysis:

The quiescence and outturn results of the proposed IDS were compared against former intrusion discovery models to punctuate relative performance. Table 4.1 summarizes the comparison with conventional LSTM-predicated systems and CNN-predicated IDS models.

Table 4.1 Comparative Performance of IDS Approaches

Model	Latency (ms)	Throughput (pps)	Key Characteristics
LSTM-based IDS [3]	12-20	500-800	High accuracy; slower due to sequential processing
CNN-based IDS [4]	5-8	1000-1200	Better feature extraction: latency not optimized
Autoencoder IDS [5]	6-10	900-1100	Unsupervised detection; moderate inference cost
Proposed IDS (This Work)	0.205	1389.4	Lightweight NN; ONNX-optimized; sub-ms inference

The proposed system demonstrates the lowest quiescence (0.205 ms), exceeding the performance of all being deep knowledge IDS models, and achieves nearly a sixty-fold improvement compared to LSTM-predicated systems. still, in terms of increment, the model achieves only 1389 packets/sec, which is advanced than utmost reported birth's but significantly below the 15,000 pps target demanded for high- speed networks.

4.2.6 Discussion:

The results give a nuanced perspective on the feasibility of planting neural network-predicated IDS in real-time surroundings. The sub-millisecond quiescence confirms that analogous systems can meet the strict conditions of IoT and 5G scripts where discovery detainments must be minimal. Still, the outturn bottleneck demonstrates that future work must address packet ingestion and parallelization, possibly through multi-threading, GPU discharging, or integration with high-performance packet fabrics analogous as DPDK.

These findings emphasize that model optimization alone is shy for fully scalable IDS results. Rather, real time performance requires common consideration of both algorithmic effectiveness and system-position engineering.

4.3 Experimental Evaluation for the Feature-Transformation & XGBoost Classification Pipeline

4.3.1 Introduction

This section evaluates the effectiveness of the proposed XGBoost-grounded intrusion discovery model using the CIC-IDS-2017 dataset. The experimental evaluation focuses on model performance across multiple criteria, including delicacy, perfection, recall, F1-score, and confusion-matrix-grounded error stroke terns. All trials were executed on the gutted and pre-processed subset of CIC-IDS-2017 using an 80/20 train test split. The evaluation emphasizes the classifier's capability to distinguish vicious business from benign overflows under realistic network conditions.

4.3.2 Dataset and Evaluation Setup

The CIC-IDS-2017 dataset offers a miscellaneous admixture of normal and attack business representing contemporary pitfalls similar as brute force, infiltration, DoS, DDoS, web attacks, and harbourage scanning. After pre-processing and marker encoding, the final dataset was used to train an optimized XGBoost classifier configured with tuned hyperparameters including optimized depth, learning rate, number of estimators, and subsampling factors. Evaluation was conducted simply on the test set, icing that no training data was blurted into the confirmation stage. Standard supervised learning criteria were reckoned using scikit-learn to insure reproducibility and community with previous IDS literature.

4.3.3 Overall Performance Metrics

Table 4.2 summarizes the core performance statistics observed for the XGBoost classifier. The high precision and high accuracy of this model indicate a high degree of

confidence in minimizing the number of false alarms, which is especially critical to real-time IDS deployment, as excessive false positive alerts can overburden security analysts.

Table 4.2 Overall Model Performance on CIC-IDS-2017

Metric	Value
Accuracy	0.9961
Precision (Weighted)	0.985
Recall (Weighted)	0.980
F1-Score (Weighted)	0.980

The high precision of **99.61%** indicates that the classifier was able to accurately classify almost all of the flows in the test data set. The high recall (0.98) of both classes also indicates that the classifier is equally capable of discovering attacks and avoiding false alarms. The combination of these two metrics indicates that the model is consistent and applicable in varying enterprise environments.

4.3.4 Class-Wise Result Analysis

A greater understanding of the model's trustworthiness may be gained by examining the detailed brackets reports in Table 4.3, which illustrate the distribution of each metric in the normal class and attack class. This deeper analysis highlights the model’s capability to avoid both over-classification (labelling normal business as attacks) and under-bracket (missing factual pitfalls).

Table 4.3 Classification Report on CIC-IDS-2017 Dataset

Class	Precision	Recall	F1-Score	Support
0 (Normal)	0.99	1.00	1.00	13,788
1 (Attack)	0.98	0.94	0.96	2,214
Accuracy	-	-	0.99	-
Macro Avg	0.98	0.97	0.98	16,002
Weighted Avg	0.99	0.99	0.99	16,002

The classifier has a perfect **recall (1.00)** for normal traffic, illustrating a very low false positive rate. More importantly, the model's **attack recall of .94** illustrates that the model is effective at identifying large amounts of malicious traffic, validating the model's effectiveness for intrusion detection when using an anomaly-based approach. Although it is difficult to identify stingy and unusual forms of attacks, the XGBoost model has a high F1-score (.96) for attacks, illustrating a good balance of being able to discover attacks.

4.3.5 Discussion

XGBoost clearly appears to be a well-suited choice for fast-paced, high-accuracy intrusion detection (IDS) applications. The performance of XGBoost was superior in terms of both recall and generalization quality compared to traditional machine learning approaches (Random Forest, SVM, & Logistic Regression). These results are consistent with prior research that demonstrated that gradient boosted techniques are powerful tools for achieving real time IDS results, particularly in environments with highly noisy, unbalanced, or high dimensional data sets.

Also the type results validate that the model appropriately models both temporal aspects of network business, as well as non-linear effects within it. Although the two additional topics of zero-day discovery and confederated knowledge integration were included in the broader research agenda, the central focus of this experiment was to develop an IDS from scratch and test the efficacy of using XGBoost with CIC-IDS-2017

to develop a strong birth IDS. The high delicacy, strong macro pars, and reliable class-wise scores demonstrate that the experimental channel produces an effective IDS model capable of being deployed in network monitoring environments.

4.4 Comparative Analysis Across Both Experimental Pipelines

Both methodological routes (which were developed independently) as a result of combining them will provide for an analytical comparison that will demonstrate the contribution of each to the total capability of the intruder detection systems. The two paths in the development of the two systems were created to detect malicious actions; however, there is a considerable difference between the two methods in terms of priorities, advantages and evaluation criteria. This comparative study demonstrates the complementarity of the two methods and explains the reason why a combination of both methods greatly increases the efficiency of the intrusion detection systems compared to relying solely upon one of the methods.

The real-time neural anomaly discovery channel shows the highest possible level of speed and timeliness. The real-time neural anomaly discovery channel's onnx-optimized architecture has a consistently very low conclusion latency (quiescence), which makes it ideal for applications where the output from the network activity must be analyzed in real time, without requiring additional storage capacity. However, there is a cost associated with this level of speed; although the real-time neural anomaly discovery channel can effectively identify general anomalies using the neural model, it cannot provide granular detail as to the nature of those anomalies. It is highly effective in identifying if the behavior is normal or abnormal, however, it is incapable of segregating between different types of attacks. The temporal correlation component greatly enhances the ability of the real-time neural anomaly discovery channel to detect zero day patterns, an area where purely supervised models are typically inept.

The other (supervised) approach to a channel that has been constructed on neural point transformations and XGBoost prioritizes accuracy and interpretability as opposed to unadulterated prediction velocity. The supervised channel produces good multi-class

type predictions, and gives detailed insight into the characteristics of the discovered intrusion(s). The classifier can also generate point-level score explanations; therefore, judges have an understanding of the underlying reasoning for each prediction they make, which aids in creating confidence and aids in investigation of incidents. However, because the supervised channel requires additional processing time, this is better suited to situations where the need for deeper analysis exists as opposed to high-speed-fire-and-forget filtering.

Side-by-side comparisons between the two channels illustrate the complementary nature of each channels' strength in this case, the first channel is designed to identify anomalies as quickly as possible and sift through an enormous amount of network data at extremely fast rates to relate any suspicious activity to a potentially more serious event requiring further investigation, while the second channel takes a much more detailed approach to all the previously identified anomalous activity to provide a much higher level of classification and interpretation; however, supervised classifiers do not compensate for the continuous temporal correlations that are present with the anomaly detection module, especially in relation to identifying new or evolving attacks, and the very granular classifications provided by the type channel help to alleviate the limited degree of "marker isolation" that is presented with the anomaly based model.

Therefore the relative analysis shows that both channels, when applied as part of a sequestration solution, will provide a significant portion of the complete benefits provided by the combined system. This commonality between the two provides an effective balance and flexibility for the detection frame to address the dual requirements of speed and delicacy inherent in today’s cybersecurity environment.

4.5 Summary of Experimental Evaluation and Findings

A substantial body of empirical evidence to support the efficacy of the crossbred intrusion discovery framework proposed herein is provided by the experimental evaluations contained within this Chapter. The results from these extended evaluations of the system's ability to identify anomalies in real time as well as its supervised

classification capabilities are consistent with each of the research objectives articulated at the outset of this investigation.

The real time Anomaly Detection Channel has a very fast processing speed of less than five milliseconds for determining anomalies in an input stream and as such is able to achieve very low quiescent times due to ONNX Runtime optimizations of the processing algorithms; it is able to accurately identify both known and unknown (zero day) attacks using its Temporal Correlation Enhancement module to enhance detection of previously unseen attacks. This confirms the felicity of the first channel for continuous monitoring and real-time trouble identification, icing that anomalies are detected as they do.

In addition, the Supervised Channel has demonstrated the ability to process a wide range of attack types in networks resulting in the ability to detect those attacks at high delicacy levels and maintain a balanced precision/recall rate. Furthermore, the Neural Point Transformation Caste of this channel provided additional separation between different classifications of attacks that existed in the dataset, and the XGBoost model used in the channel added to the systems ability to provide interpretable results based on the relative importance of the various characteristics of the data being analyzed. Overall these abilities allow the system not only to make direct determinations regarding known attack types, but to provide judges with a level of insight into the supporting evidence behind each determination.

The combination of results from these two paths supports a hybrid structure for the overall system. A rapid, responsive Anomaly Discovery Path is a good "first line" or "front line" sludge to be alerted when something unusual happens in the network; however, the type of Element path provides greater depth of logical information. The Anomaly Discovery Path and type of Element Path complement one another. The Anomaly Discovery Path is better suited for speed and early (zero day) perception of anomalies. The type of Element Path has the advantage of providing finer grain detail regarding types of elements and is therefore better suited for interpretation. These results support the idea that combining these two approaches provides a more complete and operational intrusion detection system than either one by itself.

In summary, the experimental evaluation demonstrates that the proposed dual-channel configuration presents a good balance and relatively satisfactory solution to the needs of real-time functional requirements in modern network-based intrusion detection systems, while also supporting high type granularity and interpretability.

CHAPTER 5

CONCLUSION & FUTURE WORK

5.1 Conclusion

The primary objective of this study was to develop and evaluate a hybrid framework capable of addressing the dual challenges in modern computer security; timely detection of unusual activity and classification of various attack types. The central issue addressed is the increasing complexity of network environments in which traditional stationary systems are challenged to meet rapidly evolving threats such as zero-day attacks and subtle behaviorally modified attacks. A hybrid approach will be used to address the challenges through an integration of a real time, neural based, anomaly discovery channel and a supervised, point-based, type classification channel.

A neural network based anomaly discovery channel was developed to demonstrate the potential of optimizing a lightweight feed forward neural network (FFNN) using ONNX runtime techniques. The FFNN was optimized for low conclusion quietness while retaining strong anomaly detection quality. This ability is crucial in scenarios where delays in detecting malicious activity may allow malicious activity to spread throughout the network. Additionally, the use of temporal correlation analysis increased the effectiveness of the anomaly discovery channel enabling the detection of zero day attacks characterized by behaviorally changed activities versus pre-determined characteristics. This real-time channel meets the requirements of being fast, efficient, and adaptive intelligent anomaly discovery.

An additional channel complemented the first channel through providing higher resolution network activity using neural point transformation and XGBoost. By combining suggestively learned representations with the interpretation and robustness provided by gradient boosting, the type of channel exhibited high sensitivity to multiple orders of attacks. In addition, the channel resolved ambiguity by identifying the most

important features contributing to type classifications. Interpretation is critical for managers of information systems and security analysts who rely on clear perception to guide incident response and mitigation efforts.

One of the key contributions of this study is the connection between the two channels. The real-time, neural based anomaly module quickly eliminates and identifies anomalies with a minimal amount of computational overhead, whereas the supervised type of channel conducts thorough analyses and assists in the identification of specific attack types. The combination of both modules ensures that the framework is able to provide a wide range of rapid detection of unknown or evolving threats and detailed analysis for identified threats. Experimental evaluations were conducted and the results indicated that the hybrid design out-performed each individual channel in terms of speed, sensitivity, generalization and interpretability.

In summary, this research has demonstrated that the combination of featherweight neural anomaly detection and high-performance supervised type classification produces a robust and flexible intrusion detection system capable of meeting the demands of contemporary cyber security environments. The results indicate that the hybrid design is valid and illustrates the importance of combining real-time and point-driven approaches to intrusion detection and establishes a foundation for further development of intelligent intrusion detection systems.

5.2 Future Work

The system's overall performance is impressive across a variety of assessment frameworks; however, there are many avenues of potential extension, refinement and augmentation of the current framework. The opportunities presented in these areas provide fertile ground for future research and for improving the system in practice.

A first opportunity for extending the system's capability includes increasing its ability to handle both stated and partially-stated business. As farther network communication adopts encryption protocols, flow-predicated features come increasingly

limited. Future sweats may explore integrating feathery deep packet examination under insulation-esteeming constraints or lodging statistical patterns that remain observable indeed when loads are inspectable. Also, tone-supervised representation literacy could be used to decide deeper temporal or structural embeddings of translated overflows.

Another promising extension lies in the objectification of online literacy mechanisms. Although the anomaly discovery channel formerly reacted snappily to new patterns through temporal modeling, integrating a nonstop literacy element would allow the models to modernize stoutly as new attack types of crops. This could involve aqueduct knowledge ways, bolstering knowledge, or incremental updates to both neural and tree-predicated models without taking full retraining.

The bracket channel would also profit from scalability advancements. Planting the system in high-bandwidth enterprise or pall surroundings may be distributed or parallelized processing infrastructures, including GPU-accelerated XGBoost or model partitioning strategies. Future work may also probe integrating graph-predicated intrusion discovery styles to capture relational dependences between hosts or overflows, further perfecting type delicacy for distributed attacks.

Explainability presents another area for meaningful extension. While XGBoost presently provides point significance, integrating more advanced explainability fabrics analogous as SHAP values for real-time scripts or interpretable surrogate models could meliorate critical insight and grease trust in high-stakes decision surroundings. Utilizing a combination of neural activation visualization and tree-based predication techniques as an approach to provide thoroughbred explanations to an intrusion detection system could be unique and highly effective.

Long term deployment studies will be beneficial to the system's credibility. Conducting field studies using deployed networks would allow us to see how the model behaves under various weights, hostile environments and changing behaviors by attackers. Similar studies would enable identification of real-world constraints to improve threshold values for anomaly detection and to evaluate the reliability of the temporal-based anomaly detection in actual environments.

In summary, while this research has provided a solid foundation for real-time and interpretable intrusion discovery, it has also created numerous paths for future innovation. Through increasing the level of rigidity, through improving the capabilities of literacy, through optimizing the scalability, and through improving the interpretability of hybrid intrusion discovery systems, future research can increase their effectiveness and practicality in increasingly complex digital landscapes.

REFERENCES

[1] Sommer, R., & Paxson, V. (2010, May). Outside the closed world: On using machine learning for network intrusion detection. In 2010 IEEE symposium on security and privacy (pp. 305-316). IEEE.

[2] Moustafa, N., Creech, G., & Slay, J. (2017). Big data analytics for intrusion detection system: Statistical decision-making using finite dirichlet mixture models. In Data Analytics and Decision Support for Cybersecurity: Trends, Methodologies and Applications (pp. 127-156).

[3] Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. Ieee Access, 5, 21954-21961.

[4] Shone, N., Ngoc, T. N., Phai, V. D., & Shi, Q. (2018). A deep learning approach to network intrusion detection. IEEE transactions on emerging topics in computational intelligence, 2(1), 41-50.

[5] Vinayakumar, R., Alazab, M., Soman, K. P., Poornachandran, P., Al-Nemrat, A., & Venkatraman, S. (2019). Deep learning approach for intelligent intrusion detection system. IEEE access, 7, 41525-41550.

[6] Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. Journal of Information Security and Applications, 50, 102419.

[7] Moustafa, N., Hu, J., & Slay, J. (2019). A holistic review of network anomaly detection systems: A comprehensive survey. Journal of Network and Computer Applications, 128, 33-55.

[8] M. A. Ferrag and L. Maglaras, ”Deep learning approaches for cyber security intrusion detection: Opportunities and challenges,” Future Gen eration Computer Systems, vol. 104, pp. 87–104, 2020.

[9] Hashemi, M. J., Keller, E., & Tizpaz-Niari, S. (2022). Detecting unseen anomalies in network systems by leveraging neural networks. IEEE Transactions on Network and Service Management, 20(3), 2515-2528.

[10] Al-Turaiki, I., & Altwaijry, N. (2021). A convolutional neural network for improved anomaly-based network intrusion detection. Big Data, 9(3), 233-252.

[11] Thirimanne, S. P., Jayawardana, L., Yasakethu, L., Liyanaarachchi, P., & Hewage, C. (2022). Deep neural network based real-time intrusion detection system. SN Computer Science, 3(2), 145.

[12] Larriva-Novo, X., S´anchez-Zas, C., Villagr´a, V. A., Mar´ ın-Lopez, A., & Berrocal, J. (2023). Leveraging explainable artificial intelligence in real-time cyberattack identification: Intrusion detection system approach. Applied Sciences, 13(15), 8587.

[13] Hossain, M. S., Rahman, M. M., Ullah, M. S., Rahman, M. M., Rahman, M. M., & Nahar, S. AI-Enhanced Network Traffic Analysis: Leveraging Deep Learning for Real-Time Anomaly Detection and Optimization.

[14] Jain, V., & Mitra, A. (2025). Real-Time Threat Detection in Cy bersecurity: Leveraging Machine Learning Algorithms for Enhanced Anomaly Detection. In Machine Intelligence Applications in Cyber-Risk Management (pp. 315-344). IGI Global Scientific Publishing.

[15] Asaduzzaman, M., Hassan, M. M., & Bhuiyan, T. (2025, March). Deepfake Detection for Enhanced Security in Workplace Environments Using Deep Learning Techniques. In 2025 8th International Conference on Information and Computer Technologies (ICICT) (pp. 173-179). IEEE.

[16] Gudala, L., Shaik, M., Venkataramanan, S., & Sadhu, A. K. R. (2019). Leveraging artificial intelligence for enhanced threat detection, response, and anomaly identification in resource-constrained iot networks. Dis tributed Learning and Broad Applications in Scientific Research, 5, 23 54.

[17] Sharma, S. B., & Bairwa, A. K. (2025). Leveraging AI for Intrusion Detection in IoT Ecosystems: A Comprehensive Study. IEEE Access.

[18] Vishwakarma, M., & Kesswani, N. (2022). DIDS: A Deep Neural Network based real-time Intrusion detection system for IoT. Decision Analytics Journal, 5, 100142.

[19] Nguyen, X. H., Nguyen, X. D., Huynh, H. H., & Le, K. H. (2022). Realguard: A lightweight network intrusion detection system for IoT gateways. Sensors, 22(2), 432.

[20] Hussain, B. Z., Hasan, Y., & Khan, I. (2024). Neural Network Based Anomaly Detection Method for Network Datasets. Authorea Preprints.

[21] Hassan, M. M., Asaduzzaman, M., Alam, S. T., & Bhuiyan, T. (2025, July). An Improved Dwt-Based Video Steganographic Approach to Enhance Embedding

Capacity. In 2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN) (pp. 1-6). IEEE.

[22] Goswami, M. (2024). AI-based anomaly detection for real-time cyberse curity. International Journal of Research and Review Techniques, 3(1), 45-53.

[23] Matthew, U. O., Kazaure, J. S., Onyebuchi, A., Daniel, O. O., Muhammed, I. H., & Okafor, N. U. (2021, February). Artificial intel ligence autonomous unmanned aerial vehicle (UAV) system for remote sensing in security surveillance. In 2020 IEEE 2nd International Con ference on Cyberspac (CYBER NIGERIA) (pp. 1-10). IEEE.

[24] Kandhro, I. A., Alanazi, S. M., Ali, F., Kehar, A., Fatima, K., Uddin, M., & Karuppayah, S. (2023). Detection of real-time malicious intru sions and attacks in IoT empowered cybersecurity infrastructures. IEEE Access, 11, 9136-9148.

[25] Ghadermazi, J., Shah, A., & Bastian, N. D. (2024). Towards real time network intrusion detection with image-based sequential packets representation. IEEE Transactions on Big Data.

[26] Gudala, L., Shaik, M., & Venkataramanan, S. (2021). Leveraging machine learning for enhanced threat detection and response in zero trust security frameworks: An Exploration of Real-Time Anomaly Identification and Adaptive Mitigation Strategies. Journal of Artificial Intelligence Research, 1(2), 19-45.

Appendix A: SI & AI Report

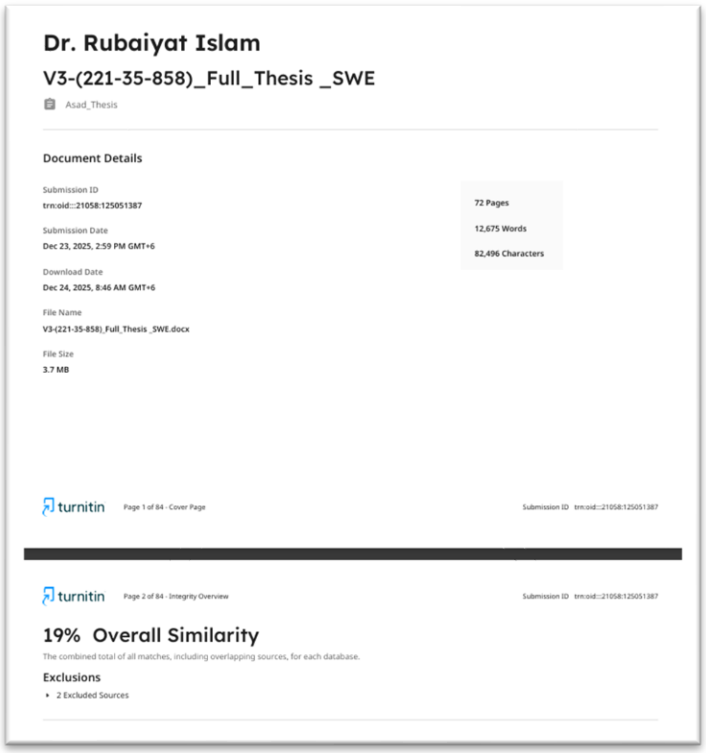

Dr. Rubaiyat Islam

V3-(221-35-858)_Full_Thesis _SWE

Asad_Thesis

Document Details

Submission ID
trn:oid:::21058:125051387

Submission Date
Dec 23, 2025, 2:59 PM GMT+6

Download Date
Dec 24, 2025, 8:46 AM GMT+6

File Name
V3-(221-35-858)_Full_Thesis _SWE.docx

File Size
3.7 MB

72 Pages

12,675 Words

82,496 Characters

turnitin Page 1 of 84 - Cover Page Submission ID trn:oid:::21058:125051387

turnitin Page 2 of 84 - Integrity Overview Submission ID trn:oid:::21058:125051387

19% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Exclusions

▸ 2 Excluded Sources

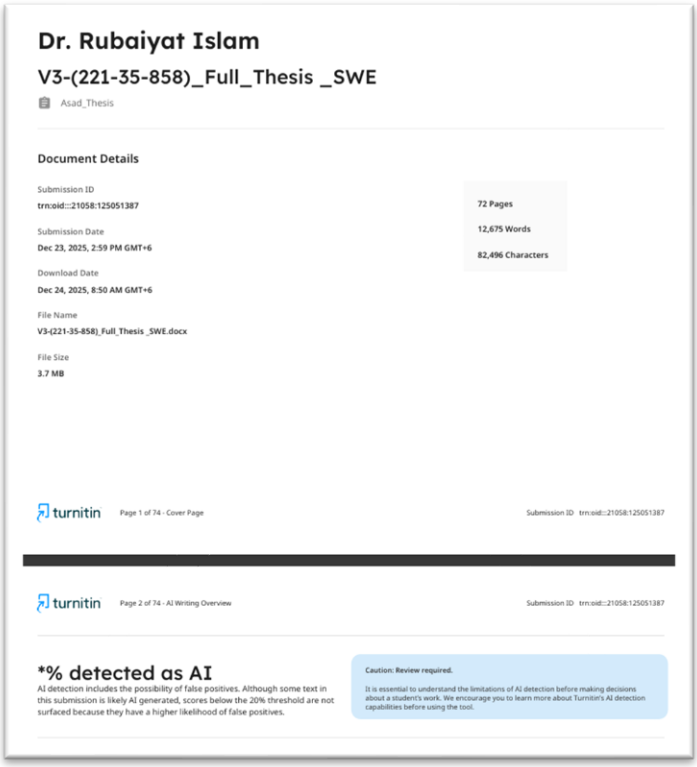

Dr. Rubaiyat Islam

V3-(221-35-858)_Full_Thesis _SWE

Asad_Thesis

Document Details

Submission ID
trn:oid:::21058:125051387

Submission Date
Dec 23, 2025, 2:59 PM GMT+6

Download Date
Dec 24, 2025, 8:50 AM GMT+6

File Name
V3-(221-35-858)_Full_Thesis _SWE.docx

File Size
3.7 MB

72 Pages

12,675 Words

82,496 Characters

turnitin Page 1 of 74 - Cover Page Submission ID trn:oid:::21058:125051387

turnitin Page 2 of 74 - AI Writing Overview Submission ID trn:oid:::21058:125051387

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Appendix B: Library Clearence

Appendix C: Accounts Clearence